Electronic Specifier design

August 2013 / Volume 3, Issue 8

www.electronicspecifier.com

If you love it, never let it go

A new generation of electronic tags could help protect our most cherished possessions

Reader Offer!
Win a PIC24F K-Series Microstick from Microchip!

In this issue:

NEED A COMPETITIVE QUOTE FAST?

TE connectivity
KEMET
Texas Instruments
molex
Panasonic
ANALOG DEVICES
muRata
BOURNS

CHALLENGE JOHN!

- Quick quotes available on 500 000 products
- 9,900 prices checked and reduced
- 7,000 lines quoted per day
- Voted best in class for delivery in UK*

*Survey based on 2,182 customers purchasing with multiple suppliers answering: "Which company do you feel provide the best delivery service?"

Submit your quote today by:

email: sales@farnell.co.uk
phone: 08447 11 11 11
website: uk.farnell.com/quotes

Farnell element14

farnell.com

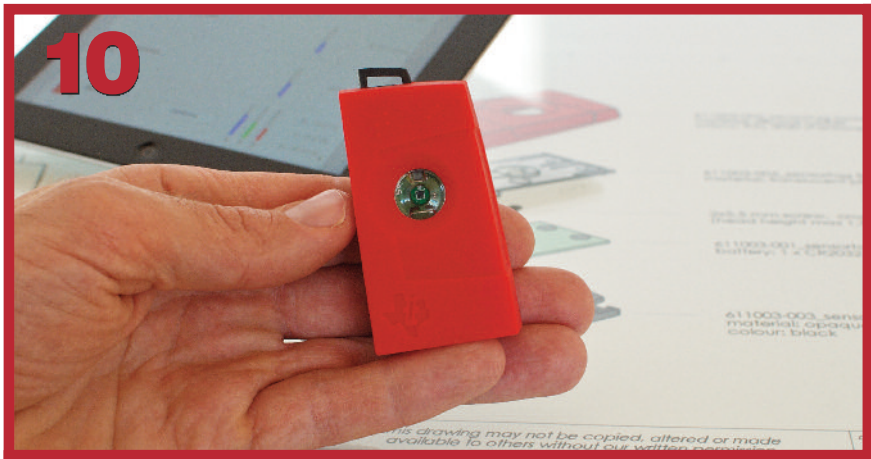

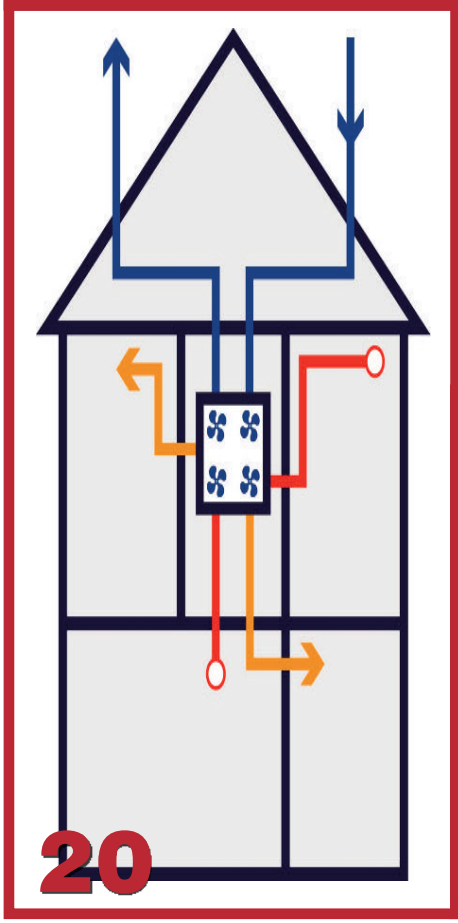

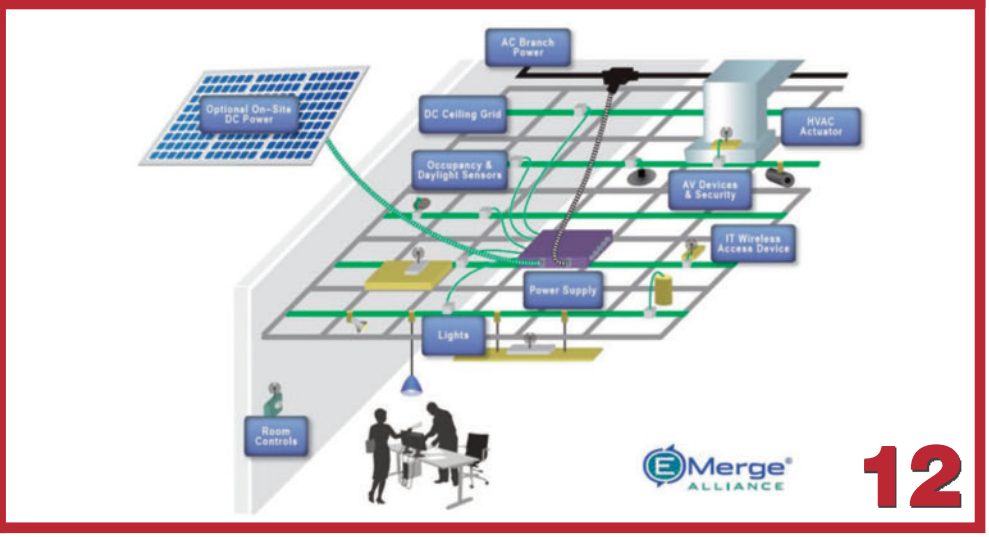

cover:

A new generation of wireless-enabled electronic tags, supported by smart phone apps, could be the next big thing in consumer must-haves

comment:

How connected are you? Not as connected as you're likely to be in the near future.

news:

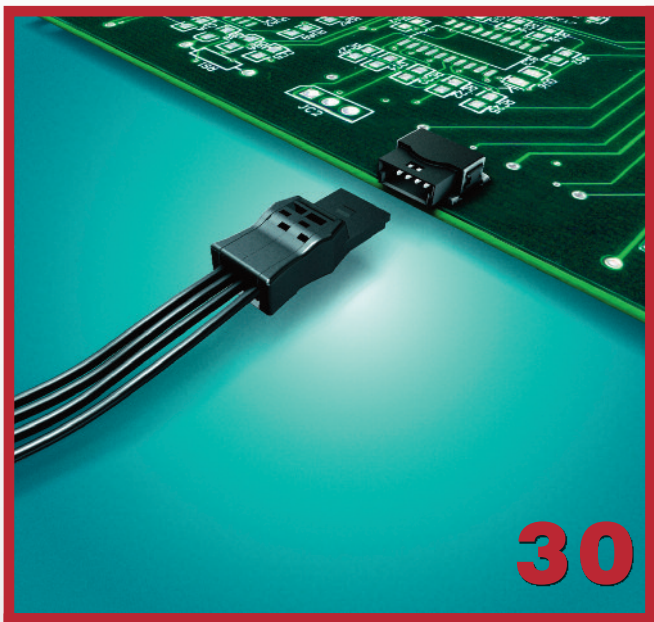

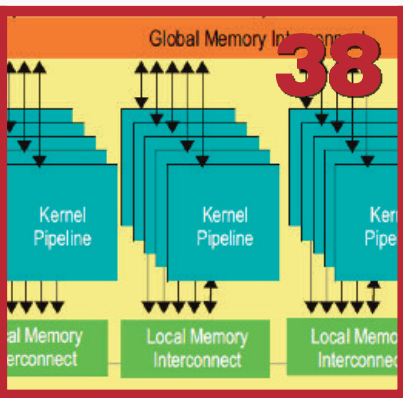

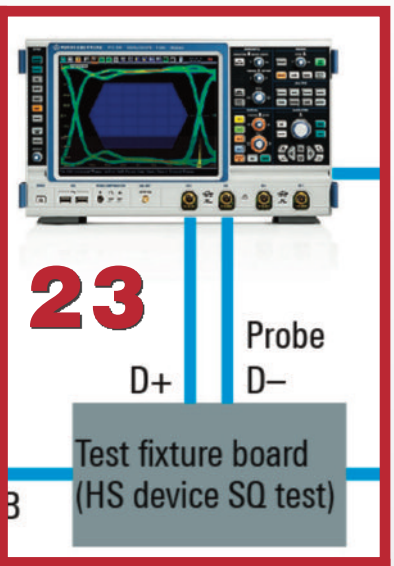

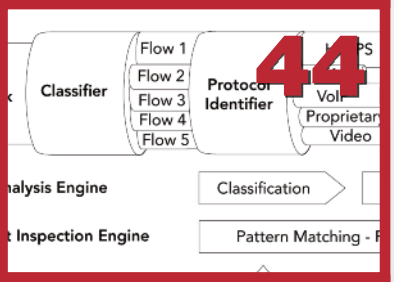

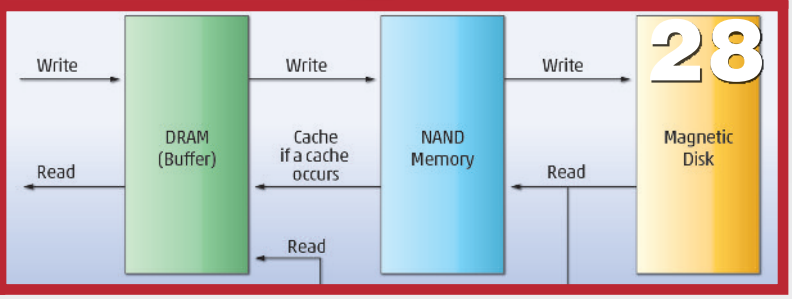

articles:

How connected are you?

The idea that, today, we're never really out of reach is supported almost entirely by our smart phones; they have the ability to relay our position from practically anywhere in the world. But despite popular perception, not everyone in the world owns a smart phone. Take this single device out of the equation and suddenly we're not as connected as we perhaps thought.

But that is most definitely changing. A new trend towards license-free wireless connectivity means that even if we aren't carrying a portable modem in the form of a smart phone, we may soon be carrying some other devices that will find a way to connect; be that through an open network powered by a common protocol, such as WiFi, or through 'piggy-backing' off someone else's smart phone.

Connectivity – in all its forms – remains key to the development of society, the prevalence of wireless connectivity may merely serve to accelerate that development.

Philip Ling, Editor

Electronic Specifier

ElectronicSpecifier Ltd.
Comice Place, Woodfalls Farm, Gravelly Ways, Laddingford, Kent ME18 6DA UK

Editor: Philip Ling
Email: phil.ling@electronicspecifier.com
Tel: +44(0)1622 679 221

Contributing Editors:
Sally Ward-Foxton Vanessa Knivett Caroline Hayes

Advertising:
Ben Price (UK Sales Manager)
Email ben.price@electronicspecifier.com
Tel: +44(0)1622 872 184 Cell: +44(0)7585 906 059

Tim Anstee (RoW)
Email: tim.anstee@electronicspecifier.com
Tel: +44(0)1622 871 944 Cell: +44(0)7771 628 990

Publisher: Steve Regnier
Email: steve.regnier@electronicspecifier.com
Tel: +44(0)1622 872 204 Fax: +44(0) 1622 870 279

Subscriptions
Subscribe to the Digital Edition of ElectronicSpecifier

Distribution Blog
Find out what's happening in the world of distribution

Scaling new lows

Xilinx has revealed the first high-level details of a new architecture it believes will outlive the next process node, allowing the same platform to be implemented in planar and FinFET processes with little or no modification.

The architecture, called UltraScale and has already been taped-out on a 20nm node, has been designed to deliver ASIC-like performance and scale from monolithic to 3D ICs. According to the company, UltraScale is already using double patterning, so 'no' changes would be needed to move to a FinFET process. It will also be possible to migrate from 20nm to 16nm and beyond, says Xilinx.

The company says it has applied ASIC design techniques to improve routing and reduce signal latency, as well as improve clock skew. Vivado, the company's IDE, will play a key role in achieving up to 90% utilisation. The architecture isn't being applied to all products; Xilinx said it wouldn't provide any real benefits to low-end FPGAs such as Artix, which will remain at 28nm, and it isn't anticipated that the Zynq family will migrate to an UltraScale platform. However, Kintex and/or Virtex are being moved to UltraScale and devices are expected to be announced before the year's end, while test chips of UltraScale using a FinFET process have also been taped out and samples are expected by the end of 2014.

Control funding

A €2.6 million fund from the German Federal Ministry of Education and Research has been directed towards developing a modular battery system with individually controlled cells. The three-year IntelliBat research project comprises BMZ GmbH, Alfred Karcher GmbH, Neutron Mikroelektronik GmbH and the Institute for Power Systems of the University of Applied Sciences Nuremberg.

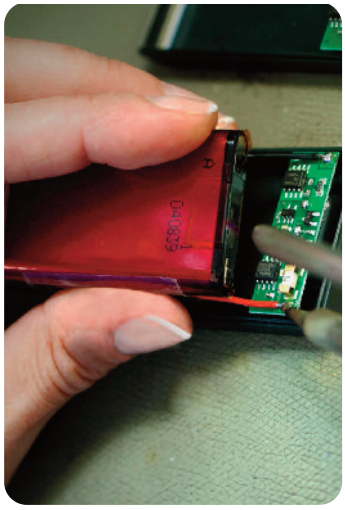

ARM validates DDC

SuVolta's Deeply Depleted Channel (DDC) technology, which is claimed to provide an easier solution than fully depleted channels but with similar benefits, has been validated by ARM, who stated that ultra-low power technologies will be vital to ensure ARM remains at the forefront of the Internet of Things.

A Cortex-M0 processor has been fabricated in a 65nm bulk planar process using DDC which, when compared to the same processor/node using conventional transistors, exhibited 50% lower total power at 350MHz and 55% increased operating speed (at matched supply voltage).

As the process is much simpler than FinFET technology, SuVolta claims it is simpler to 'insert' in to the design/fabrication flow; it also shows signs of being scaleable to as low 20nm.

SuVolta is now working with six manufacturers to bring products to market, with the first expected in early 2014.

Connecting cars

The industry's first reference design for production-line connected vehicles has been unveiled by NXP and Cohda Wireless. The MK4 integrates anti-collision technology based on the RoakLINK chipset, which provides both Vehicle-to-Vehicle (V2V) and Vehicle-to-Infrastructure (V2I) implementations, collectively referred to as V2X.

Weightless demo heralds hardware

The first public demonstration of the whitespace wireless technology, Weightless, took place recently, promising the imminent roll-out of pre-certified hardware.

The demonstration used transmission, from strain gauges mounted in a model bridge, to a Weightless terminal. The sensors showed displacement of the structure as a train passed over the bridge. The Weightless SIG now has close to 1000 members, and plans to make a series of 'significant' announcements over the next few weeks.

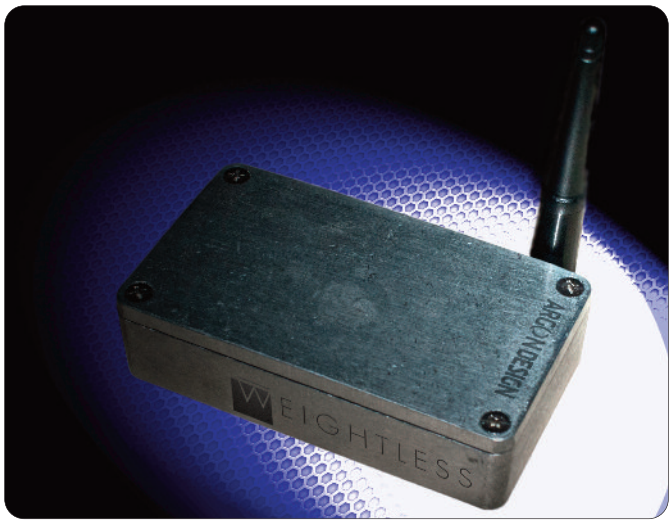

The world's first digitally enhanced power analog controller

New analog-based power management controller with integrated microcontroller

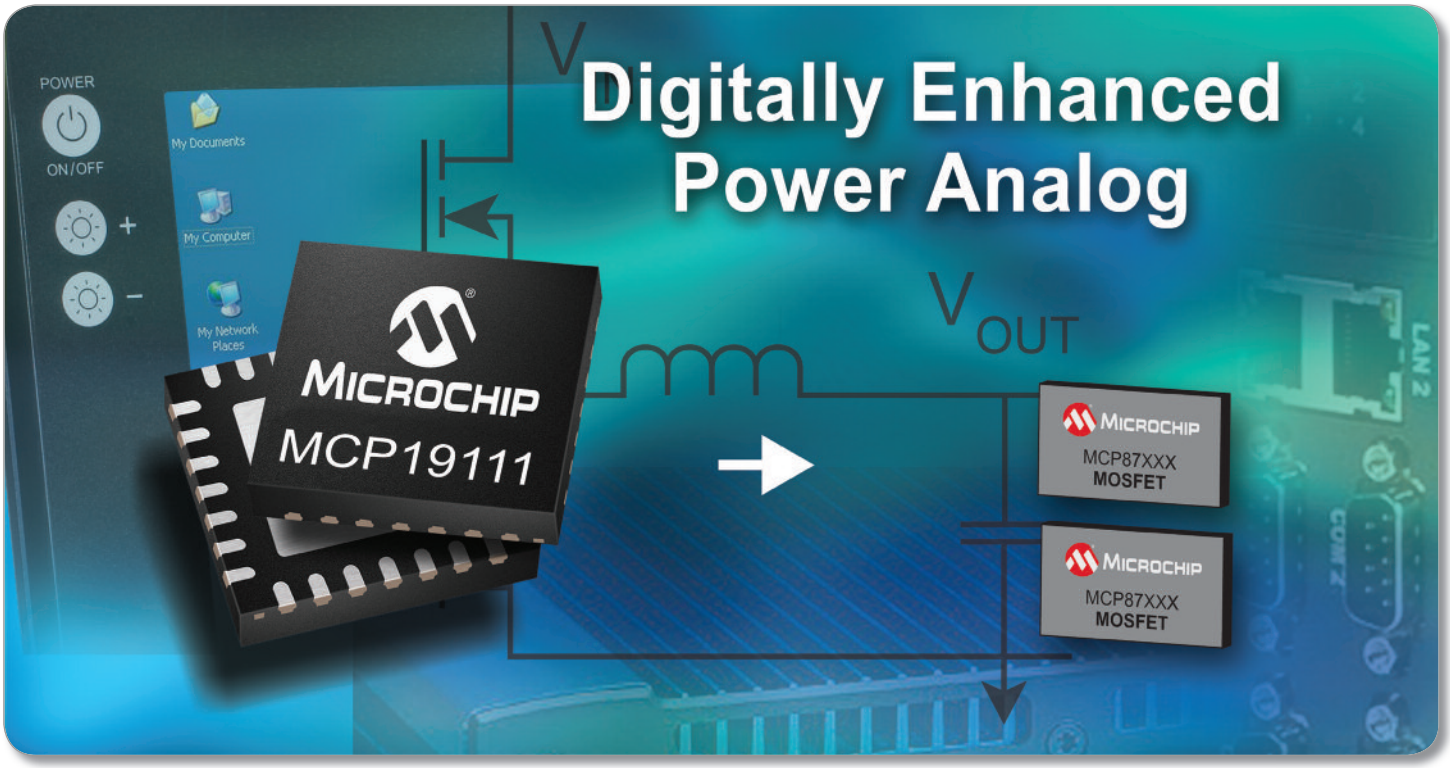

Combine the flexibility and I²C communication of digital DC/DC power conversion, with the speed, performance and resolution of analog-based control, by using Microchip's new MCP19111 power management controller.

The MCP19111 is a new hybrid, mixed-signal controller which combines analog and digital power management into a single chip. By integrating an analog-based PWM controller, a Flash-based 8-bit PIC® microcontroller, and MOSFET drivers for synchronous, step-down applications, the MCP19111 enables configurable, high-efficiency power conversion.

With transient performance of analog power conversion, the MCP19111 eliminates the need for a high-MIPS microcontroller or a high-speed A/D converter, minimising solution cost and power consumption.

To support even higher efficiency, the MCP19111 can be used to drive Microchip's latest MCP87xxx high-speed, SMPS-optimised MOSFETs. These logic-level 25V MOSFETs offer a low Figure of Merit (FOM) with on-state resistance from 1.8 mΩ to 13 mΩ to deliver higher DC/DC power conversion efficiency.

GET STARTED IN 3 EASY STEPS:

1. Evaluate the MCP19111 using the low-cost ADM00397 board
2. Combine with SMPS-optimised MCP87xxx MOSFETs
3. Customise DC/DC conversion to match your application

For more information, go to: www.microchip.com/get/eumcp19111

microchip DIRECT
www.microchipdirect.com

MICROCHIP

Microcontrollers • Digital Signal Controllers • Analog • Memory • Wireless

Low-energy servers

One of the first servers to use ARM processors has been launched by the UK's Avantek Computer, beating the likes of Dell and HP to market.

Tony Lees, MD of Avantek, said: "It's fantastic for a small company like ours to be able to beat the big boys to the market with an innovative product that's really going to point the way forward for server technology."

The new server consumes just 5W, compared to the more usual 100W of an x86-based server, said the company, while generating commensurately less heat. Based on the ARM Cortex-A9, the Calxeda EnergyCore server-on-a-chip uses solid state drives and comprises a total of 48 server nodes, each with a quad-core.

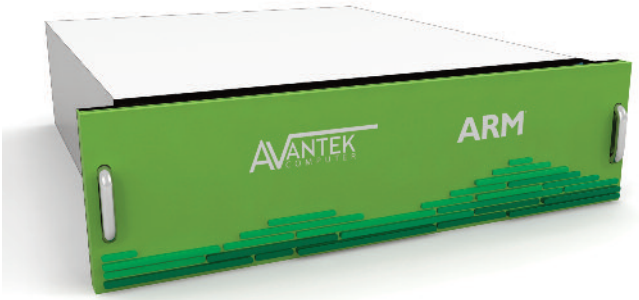

BLE for PSoC

Cypress has announced it has qualified a Bluetooth Low Energy radio which it intends to integrate in to its programmable and capacitive touch platforms.

The company intends to develop single-chip solutions combining the radio with PSoC, CapSense and TrueTouch technologies, targeting Bluetooth Smart applications such as PC peripherals, remote controls, wearable electronics and portable medical devices.

The latest release of the Android operating system now supports Bluetooth 4.0, which includes Bluetooth Smart, creating new opportunities for 'appcessories'.

Paul Williamson, Director of Low Power Wireless at CSR, said: "With Bluetooth Smart there is a single profile that can be used to communicate with any accessory; the intelligence lives in the app, not the operating system, so a new appcessory can be developed and launched as fast an an application."

Fast charging

Power Integrations has implemented Qualcomm's Quick Charge 2.0 protocol in its CHY100 IC, allowing mobile devices to charge up to 75% faster than conventional fixed-voltage chargers.

The device will enable manufacturers to add Quick Charge 2.0 to their wall chargers, which should enable smart phones with high-capacity batteries to charge in as little as one hour.

Qualcomm chose Power Integrations as its lead strategic partner to develop a production-ready IC, which it achieved in less than six months.

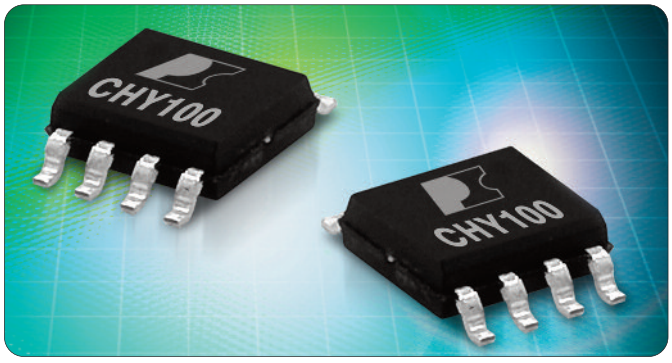

Java support for ARMv8

ARM has announced a multi-year agreement with Oracle to add support for Java Standard Edition (Java SE) to the ARMv8 64bit platforms, highlighting the growing use — and importance — of ARM architectures in server and network infrastructure.

Cool EVs

The Technology Strategic Board has awarded a grant of £317,000 to Versarien, to develop thermal management solutions for the power electronic modules of electric vehicles and plug-in hybrid electric vehicles (PHEVs).

The company is to collaborate with Applied Materials Technology and Dynex Semiconductor. Versarien intends to use its patented micro-porous metals with advanced finishing and bonding techniques to meet the low thermal budgets in EVs and PHEVs. The project is expected to take two years to complete.

Han-Eco®

People | Power | Partnership

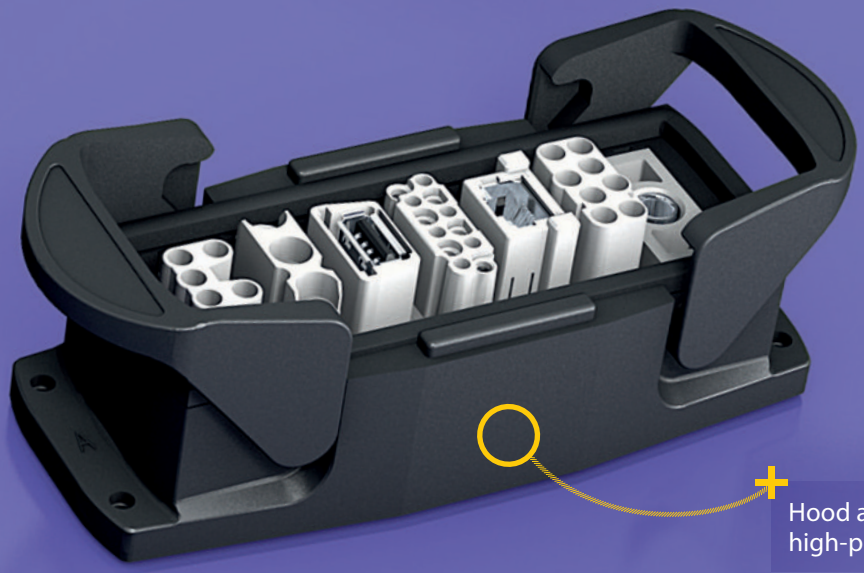

Quality with Reduced Weight

Han-Eco® – the new HARTING housing series made of robust high-performance plastic. The new Han-Eco® thermoplastic housing has excellent mechanical properties and it is lightweight. Moreover it is highly resistant against environmental influences and can additionally use one extra module from the Han-Modular® series. Versions of the hoods come with either straight or angled entries and housings, bulkhead mounting, can be supplied in all Han® B sizes. Versatile use at lower cost.
HARTING: Pushing Performance to innovative solutions.

HARTING®

Pushing Performance

If you love it, you should never let it go

With the proliferation of connectivity, consumers are being encouraged and enabled to implement their own security measures to protect the things they love most, as Philip Ling discovers.

Tagging isn't new; who didn't have name tags on their clothes and belongings at school? Who doesn't use luggage tags when travelling abroad, and how many of us now have a reward scheme associated with out house/car keys if lost? Tagging makes inherent sense, to increase the chances of retrieving lost property, but what about as a preventative measure to avoid losses in the first place?

That's the premise of the latest wave of 'smart tags' that are hitting the market, effectively enabled by a creating and maintaining a wireless connection to something that most of us already cherish; our smart phones.

Smart tags that use wireless tethering aren't new; there are numerous examples of wireless tags already on the market that create a virtual 'leash' between a phone and the tag, used predominantly for proximity sensing – if the object wanders too far from the smart device, an alarm sounds.

However, the latest generation of electronic tags take this a step further, by integrating more forms of wireless connectivity, more sensors and greater functionality.

TI's Sensor Tag development platform has already been used to bring two products to market

Hypertag

As with most things today, the Cloud will play a key role in making next-gen tags really useful. One example of this comes from Tile, a startup using crowd funding (enabled by the Selfstarter open source framework) that has developed a tag based on Bluetooth Low Energy, which aims to take tethering to a new level.

Tile has been developed with simplicity in mind; they are designed to be used for 12 months without charging or changing a battery (after a year, they're returned for recycling and customers replace it with a new, up-to-date Tile). More significantly they are intended to operate symbiotically with a smart phone via an application that, arguably, provides the real intelligence.

Essentially, any tag can be located via an smart phone running the app; it means that, even when out of range of the owner's phone, the tag can still be located as long as it is within range of

another phone running the app. The company assures that security is preserved and that any phone used as a 'stepping stone' back to the owner remains unaware that a tag is 'borrowing' its connectivity.

Obviously, the efficacy of this approach relies less on the number of tags in circulation and more on the number of people running the app on their smart phone, and so the company is promoting the solution as 'the world's largest lost and found'. Find out more at www.thetileapp.com

Makes sense

A key feature of the new tags is their increased 'intelligence' and while this may sometimes be more about the associated smart phone app, it will also be achieved through the integration of more sensors.

A good example of this comes in the form of the Wireless Sensor Tag, produced by Cao Gadgets LLC, which integrated a 3D compass, temperature and moisture sensors, along with a Hall effect sensor. According to founder Zhiheng Cao, there are also plans to integrate a humidity sensor.

Rather than using Bluetooth Low Energy operating at 2.4GHz, the Wireless Sensor Tag operates at 433MHz using a proprietary protocol and is driven by a PIC16 microcontroller, and features a replaceable battery that should power the tag for up to 3 years. Also, while there is a smart phone application for the tag, its gateway is formed by either a dedicated hardware box connected via Ethernet, or a Windows-based PC. This 'hub' network approach makes it less suited to out-of-home tagging and more for monitoring your home environment, from the temperature in your fridge to the status of your garage door. As long as you can connect to the internet, you can access the tags remotely. Visit www.wirelesstag.net for more information.

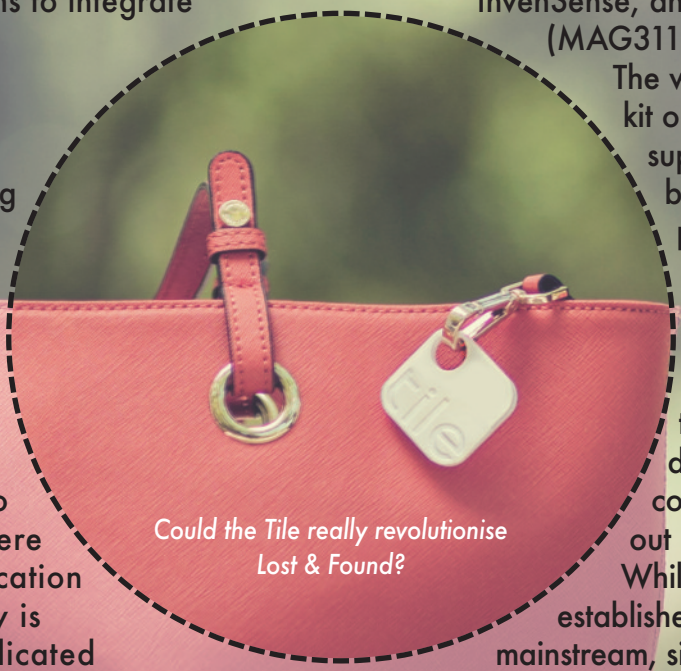

Could the Tile really revolutionise Lost & Found?

Making your own

The 'need' for wireless tags will grow as their applications evolve, but clearly they have a wide scope, perhaps even wider than the often-cited Smart Watches that are just around the corner. The ability to integrate any number of sensors, specifically MEMS-based sensors that offer low size, low power and high reliability, creates a compelling value-proposition for consumers and manufacturers. Couple that with the abundance of connected devices and the relative ease of developing a companion application, then the future for smart tags would indeed seem bright.

One solution that has already been used to bring products to market is the SensorTag development platform from Texas Instruments, based on its CC2541 Bluetooth Low Energy compliant SoC (integrated 8051 MCU and RF front-end). The kit comprises six sensors: an IR temperature sensor (TMP006) from TI; a humidity sensor (SHT21) from Sensirion; a pressure sensor (T5400) from Epcos; an accelerometer (KXTJ9) from Kionix; a gyroscope (IMU-3000) from InvenSense, and a magnetometer (MAG3110) from Freescale.

The wireless SoC at the heart of the kit operates at 2.4GHz and supports Bluetooth Low Energy, but can also support proprietary protocols.

However, the current focus is to support iOS devices and, as such, there is a companion app, which can connect to the tag and monitor its various data streams, as well as configure the tag. You can find out more at www.ti.com/sensortag

While wireless tethering is well established it hasn't exactly become mainstream, similarly NFC-enabled tags have yet to proliferate. And while the high anticipation that surrounds smart watches may, or may not, result in a new wave of consumer 'must-haves', perhaps the smart tag could just steal the limelight. Early examples indicate they could offer almost limitless scope for innovation while riding on the 'Internet of Things' trend that is widely regarded as being on the brink of explosive growth.

Return to contents page.

Was Edison Right?

A new movement in the US is seeking to bring back DC power distribution at the building level, which can help reduce power losses, as Sally Ward-Foxton discovers.

Back in the 1880s, the so-called 'War of Currents' was raging, with Thomas Edison supporting DC transmission of electrical power, opposed by George Westinghouse and various European companies, who favoured AC. At the beginning of the War of Currents, DC was standard in homes in the US, with Edison holding all the patents for DC transmission. Heavy copper wires were used for transmission but the load still had to be within a mile or so of the generator. The development of the closed-core shunt connected transformer in Hungary in 1884 gave AC the boost it needed to win the war. Stepping AC voltages up and down became easy, so transmission across many miles at high voltage with minimal losses was now possible.

Fast forward to today. The increasing electronics content in homes, businesses and cars rely on semiconductor devices, all of which require a DC power supply. AC-DC converters provide this power supply, but not without losses. The other part of the problem is that environmental and financial driving forces are increasing the percentage of power generated off-grid, typically by PV solar panels or wind turbines, which produce DC power. Currently, power from solar panels is converted from DC to AC by an inverter,

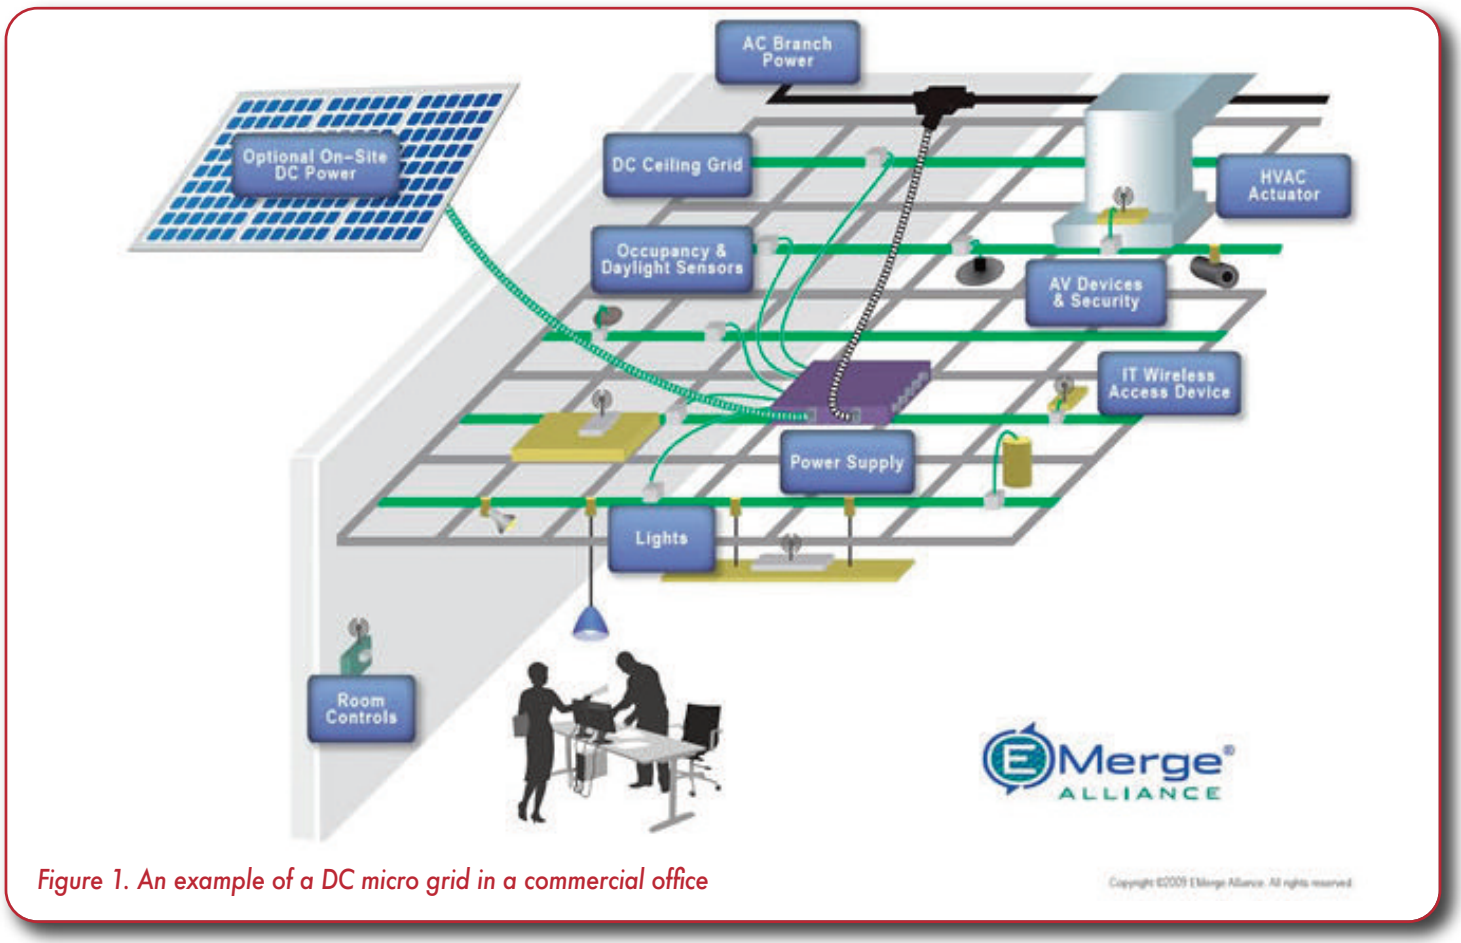

Figure 1. An example of a DC micro grid in a commercial office

Which device would you design in if you knew it would be available forever?

Design with confidence using Rochester Electronics' *Extension-of-Life®* process

Have you already received an end-of-life announcement?
Do you anticipate receiving another one?
How would you like to avoid them forever?

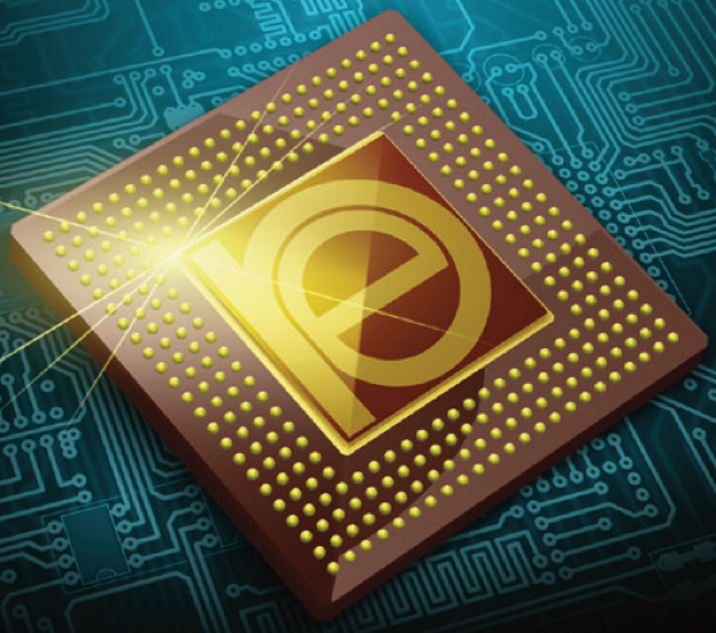

Extension-of-Life® programs from Rochester Electronics solve your end-of-life problems and keep you focused on your next development program.

For solutions go to **www.ExtensionofLife.com**

... or **SEARCH NOW!**

Click here to search for factory-direct inventory

Rochester Electronics®

Transforming Semiconductor End-of-Life into "Extension-of-Life."®

phone **978.462.9332** · **www.rocelec.com**

fed into the mains, then converted from AC back to DC by a transformer in an electronic device's power supply. There are substantial energy losses associated with both conversions.

A better way?

So, was Edison right? Is there a case for DC power distribution? Well, the laws of physics can't be changed, so DC still isn't suitable for long distance transmission. However, a movement in the US is proposing a hybrid solution, with AC used for long distance transmission and DC used in smaller power grids at the building level, called 'micro grids'. These micro grids may suit lighting ballasts, for example. The idea is to avoid having a lossy AC-DC converter on every single lamp in an office block, instead converting from AC to DC in bulk in a more efficient way and supplying all the lighting from that. An added bonus is the option to connect the micro grids directly to energy sources such as solar panels that produce DC anyway, doing away with the inverter. The resulting systems use less electronics and are less complex, adding to reliability.

The idea is gaining credibility. A new industry group, called the EMerge Alliance, has been set up to promote the use of these DC micro grids. The Alliance has also published a standard, called Occupied Space, for the use of such micro grids, as shown in Figure 1.

"Efforts to create a Smart Grid need smarter buildings that use more adaptable power infrastructures that can minimise conversion losses and make measuring and controlling power easier," says Alliance Chairman Brian Patterson. "This helps to reduce the overall load on our energy resources and make the use of alternate clean energy generation more likely... we believe that ongoing and increasing demand for improved reliability and energy efficiency across all areas of commercial buildings provides the need for our broad platform."

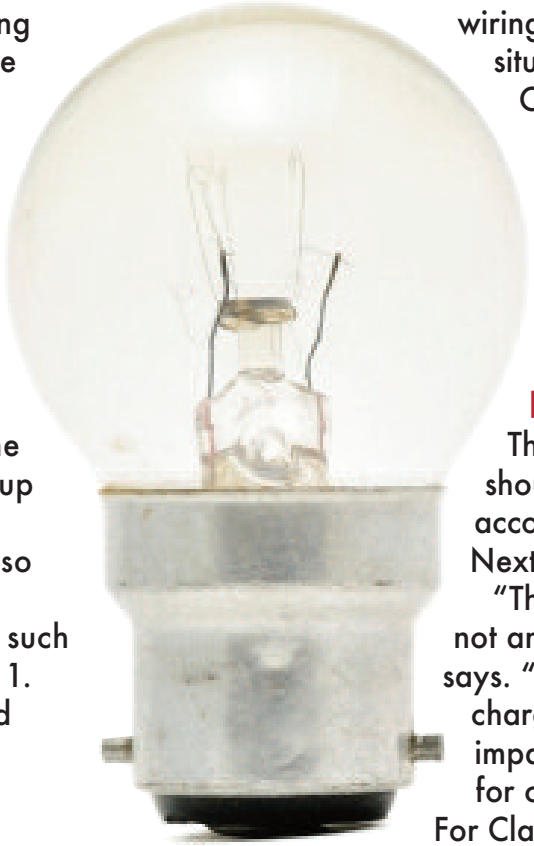

Image courtesy of ningmilo at FreeDigitalPhotos.net

Wiring losses were the main reason for DC's demise in the 1880s, so the EMerge Alliance is advocating 24V DC microgrids used over distances that won't materially affect efficiency. The length of the low voltage cables that connect power supplies to distribution busses are limited by the Occupied Space standard.

"AC-DC power supplies will be dispersed throughout a space in order to purposely keep their output cable lengths short," Patterson says. "This means that individual supply components will likely be designed to handle loads less than 100 Watts."

Distribution wiring is maintained at 12 AWG, the same as traditional AC circuits, but because this wiring is only carrying 24V, it doesn't generally require metal jacketing, metal junction boxes, ground wires or other protection means, keeping the cost of installation low. Use of existing wiring may even be possible, in some situations. Installations fall under Class 2 of the US's National Electrical Code (limited to 100 Volt-Amps), so can deliver up to 4.1A of current at 24V and still provide acceptable protection from electrical shock and fire initiation.

Designing real applications

The design of such DC microgrids should be relatively straightforward, according to Paul Savage, CEO of Nextek Power Systems.

"The first thing to know is that it's not an all-or-nothing proposition," he says. "A load like lighting or vehicle charging can be done without impacting the design considerations for conventional AC power delivery. For Class 2 systems, it's a snap, because the user-facing power is below the shock and startle hazard."

Nextek makes power server modules (PSMs) which are at the heart of a DC micro grid system, converting 240V AC power to 24V DC. The company also makes an interface module to add power from solar PV cells to the DC bus, that works in combination with the PSM. Savage says that Nextek can also supply an IPv6 wireless control node, called the Skyrouter, that controls the PSMs in a mesh network.

As far as end products go, Nextek manufactures DC-compatible ceiling fans under its Fanworks brand. Compatible lighting products are available from Philips, Osram, Acuity, Cooper and others. LED lighting is particularly well suited to the DC micro grid because it is driven electronically, but electronic fluorescent lighting can be modified to work with DC as well. Luminaires with DC input ballasts or LED drivers use the same lamps, have the same output lumens and wattage per fixture, and have the same lighting performance characteristics as AC-input devices. The efficiency gain for LED lighting installations can be 10-15%, according to the EMerge Alliance.

The future?

DC micro grids may soon be coming to homes as well as commercial buildings. Nextek is working with NextEnergy and Champion Homes on a demonstration called the NextHome, a DC connected house. NextHome will showcase companies and allow them to test, demonstrate and commercialise their technologies in a real world residential home. The project should be up and running by the end of the year.

"I think it's a global, inexorable trend that will bring benefits to the rich and the poor," says Savage, referring to the company's STAR trailers: mobile solar-powered electricity generators for rural areas of developing countries that are completely off-grid. Of course, the trailers take DC power generated by the solar panels and directly power homes and small businesses with it – no conversion to AC required. Ultimately, he says, anything with a semiconductor in it will benefit.

Meanwhile, the EMerge Alliance's vision is firmly focussed on DC micro grids throughout commercial buildings, including the occupied space, data centres, outdoor applications and building services, according to Chairman Brian Patterson.

"We see great opportunities with lighting and computing technology as well as electric vehicle charging and larger building loads such as HVAC, motor loads and high bay/industrial applications which are often DC based and are ideally suited to micro grids," Patterson says.

Return to contents page.

The smaller the better

With demand for ever-smaller portable devices, developments in the design and manufacture of miniature switches can really drive down the physical volume of consumer devices. By Jerome Smolinski

Judging by the way some users handle portable consumer electronics, they can be considered suitable for harsh environments. Cell phones, MP3 players, tablets and other portable electronic devices have become ubiquitous personal and professional tools that are used continually throughout the day, and not always with the greatest of care. As a result, switch manufacturers must create new rugged miniature switches that combine significant space and weight reductions with ruggedness and long operating lives. These miniature switches must function in the same reliable, consistent manner as the more substantially-sized industrial designs; yet without sacrificing functionality, performance or extended lifespans. Switch manufacturers that offer value-added services, including manufacturing modules and custom assemblies, are able to deliver complete electromechanical solutions that not only meet the size and performance requirements, but can also withstand the torments of vibration and shock.

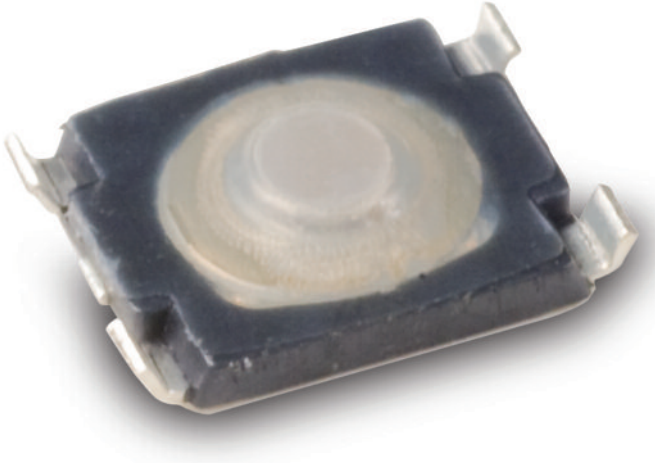

The KMT0 Series switches from C&K Components are designed with sealing Teflon films and adhesives to prevent the ingress of fluids, and provide maximum reliability

Double action switches

Low-profile double action, double-tactile switches are now being designed into handheld and portable electronic devices that require specific ergonomics, such as the focus and shutter release on a digital camera or mobile phone. These advanced double-action, double-tactile switches provide fast, repetitive action in this type of handheld consumer electronics applications. The increased functionality from the miniature double-action, double-tactile switches enables designers to eliminate components from their designs, decreasing the weight and size of their end product.

Some miniature double action, double-tactile switches available on the market today are offered in package sizes as small as 3.7mm x 3.0mm x 0.95mm, including the height of the actuator. Standard pushbutton switches have a typical life cycle of around 30,000 actuations, while typical tactile switches can reach a range of 100,000 and 600,000 actuations; with a target goal of up to 1,000,000 for future tactile switch designs. This long operating life is due in part to a rated resistance to shock and vibration of 10 to 500Hz and an operating temperature of -40˚C to +85˚C.

Pico switches

Ultra-miniaturised pico switches improve upon traditional switch designs. Unlike detect switches or pushbuttons, ultra-miniaturised top-actuated pico switches are smaller and more flexible by design. The KXT Series from C&K for example, employs small domes between one and two

Stay in communication, wherever your location, with extra battery power from Powersolve portable power banks and solar chargers!

Features

- High quality 3.7V lithium polymer batteries from 2800mAh to 10400mA.
- Considerably extends the battery life of any mobile device
- Fully protected by intelligent charging and protection circuits
- Charge any mobile device that can be charged via standard USB output, such as smart phones,tablet PC's, iPads, cameras, etc.
- 5V 1A or 2A USB outputs
- Can be charged from a standard 5V 1A or 2A USB charger or PC USB output
- Solar versions can be charged by USB charger or sunlight
- Ultra slim compact design fits easily into a pocket, briefcase or handbag
- 12 months warranty
- Ideal for corporate gifts or to keep your salesforce mobile

POWER SOLVE

www.powersolve.co.uk
Tel: 44-1635-521858
sales@powersolve.co.uk

millimetres within new types of switch structures, representing significant size reductions over nano switches. These truly miniature switches measure 2.2mm x 2.0mm x 0.5mm. A longer version at 2.8mm are also available. Ultra-miniaturised top-actuated pico switches usually provide multiple actuation forces between 1.0N and 2.2N with a 30% tactile minimum. Pico switches also typically provide an operating life of 300,000 cycles with an IP68-rated sealing. The ultra-miniaturised top-actuated pico switches deliver an extended lifespan and rugged design through the use of advanced films and strong redesigned bases, making them an ideal choice for designers working with limited space, but who cannot afford to sacrifice switch or signal quality.

Ultra-miniature pico switches have become increasingly popular for behind-the-ear hearing aids, where all components must be small in size, yet function consistently and effectively. The ultra-miniature top-actuated pico switch helps designers reduce the overall size of the hearing aid, thereby increasing user comfort and reducing the visibility of the device behind the ear.

Pedometers also benefit from switch miniaturisation through a reduction in the overall weight and assembly of the wrist portable version. The pedometer have to work under harsh conditions and keep working for a high number of actuations, making reliability the top concern when considering a new switch design.

Tactile switches are inherently smaller and more flexible than many other switch technologies. These types of switches afford the capability of multiple mounting and actuation configurations, providing greater flexibility along with the ability to configure the switch for application-specific needs. Typical tactile switches can reach a range of 100,000 and 600,000 actuations, respectively, with a target goal of up to 1,000,000 for future tactile switch designs. Their long operating life is due in part to a resistance to shock and vibration of 10 to 500Hz and an operating temperature of -40˚C to +85˚C.

Scroll switches are a popular solution in shrinking electronics because of their miniaturised packaging and rugged functionality. Scroll switches typically maintain longer lifespans than pushbutton switches because there are fewer aspects of the switch that can be compromised due to tactile interference. Scroll switches typically feature life cycles of more than one million central select actuations and one million rotation actuations (each direction), with a maximum contact resistance of 300mΩ and inherently strong shock and vibration protection, whereas conventional switch designs can typically only reach 200,000 cycles.

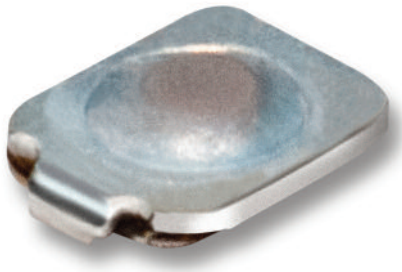

The KXT Series

Designed specifically for applications where height or surface area are restricted, sub-miniature slide switches combine a low-profile with an extended electrical and mechanical life. Available in glass-filled housings with 4/6 nylon (UL94V-2) actuators, these sub-miniature slide switches feature a stainless steel cover and return spring to allow for a 100,000-cycle standard lifespan. Measuring as small as 2.6mm by 7.6mm, the surface-mount switches are ideal for a range of applications, including on/off touch screens, handheld games, remote controls and instrumentation applications.

Side-actuated DIP switches combine miniaturisation and high-reliability for logic switching within computers. The DIP switches' combination of numerous mounting options and a low board-mounted profile increases available PCB space and offers a uniquely flexible solution for design engineers. These half-pitch DIP switches feature a dual bifurcated wiping action that ensures reliable contact. When mounted horizontally, the side-actuation mechanism provides visual confirmation of the on/off status, ideal for logic switching grids where multiple signal paths must be manually routed via switchboards.

Author profile: Jerome Smolinski is Senior Product Manager at C&K Components

These miniature DIP switches measure as little as 4.16mm above the PC board and feature a 1.28mm layout pitch. Process-sealed for surface mount wash processing, the device features gold over nickel over copper alloy contacts and a UL 94V-0 PPS flame retardant housing.

Modules and assemblies

To provide a complete solution, switch manufacturers can also design and manufacture modules and assemblies that meet customer-specific application requirements. Switch manufacturers with engineering expertise can quickly and efficiently integrate switches and electronics (from simple to complex circuitry) into attractive and functional packages that assemble neatly into customers' finished products. These value-added services can also include manufacturing inserts and injection moulded parts, stamped parts and lead frames to provide a one-stop-shop for customers.

Advanced switch manufacturers can not only refine a switch's sound (tunability), feel (haptics), and control (interface) within assemblies and modules to suit customer preferences, but also offer a wide range of additional build options, including decoration, lighting, connections and sealing. Incorporating custom graphics, logos, textures and finishes onto critical surfaces also provide a benefit to the end product. Using two-shot moulding, paint and laser etch, pad printing, and plating, these decorative solutions are designed to stand out and fit easily into a portable device, such as an iPhone or iPad. Lighting options include backlight accenting and light piping to enhance a product's appearance and practicality to a design.

Miniaturiszation and reliability are two of the more important considerations in the consumer device market today. The development of innovative miniature switches has helped enable design engineers to continue this trend. Because electrical and mechanical specifications in consumer applications are critical – tolerances are extremely tight and product profiles are small, and development cycles and lead time requirements are short – switch manufacturers that not only develop flexible and reliable devices, but also add value by manufacturing modules and assemblies, are at a distinct advantage for design wins.

More from C&K Components
Return to contents page.

Making sense of heat recovery

Mechanical ventilation and heat recovery systems are becoming increasingly necessary as we near the Zero Carbon goals for dwellings by 2016, Mike Farrer highlights the need for the highest possible efficiency power sources to perform over a wide operating window.

In general, legislation exists for good reason, in any industry. In the building industry legislation provides regulatory requirements that deliver a safer and increasingly more sustainable end-product. The drive towards sustainability isn't restricted to the building industry, of course; it now touches and influences every industry and permeates the entire supply chain.

Legislation also drives innovation; out of necessity, manufacturers who are obliged to comply with legislation must continuously strive for ways to differentiate their product or service, in what can be viewed as an increasingly homogeneous landscape. However, invariably legislation and regulation exists to protect consumers from the 'worst-case scenario'; imposing upper limitations that allow significant room for improvement. And it is here where real innovation can shine through.

Sustainability and 'going green' is a case in point. While opinions will always conflict, it's a fact that the decision from global governments to adopt 'greener' and more sustainable practices now spans design, manufacturing, distribution, and end-of-life across practically every industry at some level. The use of harmful chemicals and base materials in now highly restricted in the electronics industry, for example, while European Directives now dictate how manufacturers must also make provision for the disposal of such products.

In the building industry the focus is less on end-of-life and more on its preservation; specifically, buildings constructed today are expected to be in use for many decades and so it is their impact on the environment over their 'working life' that is now coming under closer scrutiny.

MVHR makes sense

With a target of making all new homes Zero Carbon by 2016, an important aspect of new builds is their efficiency in heating and/or cooling. Modern homes must now comply with legislation that requires them to be air tight, maximising their ability to lower heating costs through insulation and efficient windows/doors. Targets propose a roadmap of waypoints in improving insulation and heating systems, leading to the eventual Zero Carbon goal.

This isn't happening in isolation, of course; other aspects of legislation also dictate the restricted use or cessation of harmful components in building materials, coupled with a reduction in the contribution to CO2 emissions. On the face of it, an airtight building that neither harbours or vents harmful materials, or contributes any CO2 to the environment would provide the ideal solution.

But this presents its own problems, specifically in dealing with the affects of stale air and condensation. The

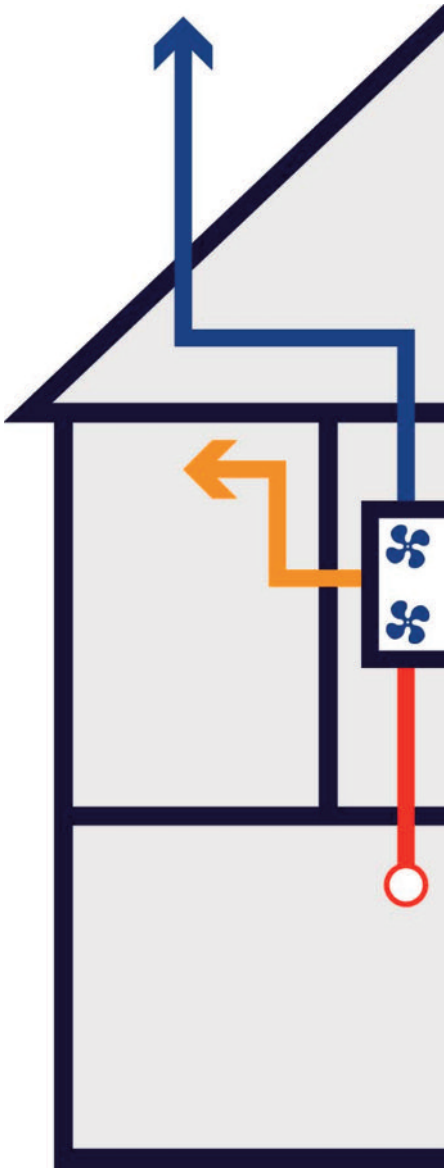

solution to this is to vent stale air to the outside and replace it with fresh air. However, while not harmful to the environment, it is clearly wasteful in terms of the heat lost by venting warm air and replacing it with cooler albeit fresher air.

Mechanical Ventilation and Heat Recovery (MVHR) helps overcome this problem, by recovering the heat from the warmer internal air and using it to heat the fresh air from outside, before it is introduced in to the property. MVHR systems are relatively simple but provide an important part of the Zero Carbon building.

In order to preserve the overall goal of reducing emissions, the performance of an MVHR system is also subject to inspection. Specifically, their performance can be submitted in accordance to the Standard Assessment Procedure Appendix Q, which makes performance data for installations of specific equipment available to energy performance assessors. It uses tests and methodologies that integrate within the applicable SAP calculation version.

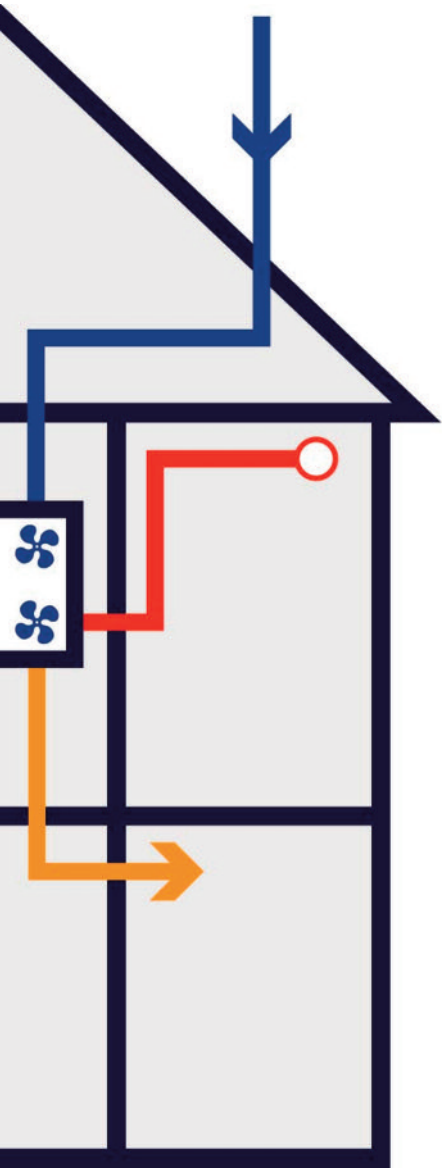

This data contributes to the overall energy performance of a dwelling and is becoming increasingly relevant as the Zero Carbon objectives of 2016 draw closer. There are now a number of MVHR equipment manufacturers and equipment listed on the Appendix Q website, allowing specifiers to select the most appropriate solution for their needs.

Driving efficiency

The performance of an MVHR installation is dependent on a number of parameters, some of which are installation-dependent – such as using rigid or flexible ducting – and some which are determined purely by the equipment and components used, like the kind of fans employed, their power supplies and control systems.

Predominantly, the components that make up the electronic subsystems are developed and assembled by a range of manufacturers, such as those listed in the database of performance data maintained by SAP (www.sap-appendixq.org).

This database of performance data describes key aspects of an MVHR system's operational parameters measured using SAP's own criteria. These include (and are largely defined by) the power needed to (re)circulate a given volume of air based on the number of (wet) rooms in the dwelling. By example, the tested airflow for a kitchen is 13 Litres/second and for wet rooms (bathrooms, for example) it's 8L/s.

The SAP Appendix Q specification measure this airflow in terms of electrical energy, or Watts. The result is a figure measured in Watts/Litre/second, referred to as the Specific Fan Power (SFP). The current specification calls for fairly modest SFP figures; <1.5 W/L/s and heat recovery efficiency >70% . Already the database of performance data lists equipment/manufacturers that achieve much better figures.

The challenge these manufacturers face, however, is improving these already impressive performance figures under the variable airflow rates required by SAP Appendix Q; something that is almost entirely dependent on the fan type used. Inevitably the requirements will get tougher.

Powering MVHR

Fans are, generally speaking, relatively cheap. However, their cost is directly related to their efficiency, which in turn is dependent on their operating method. The cheapest and simplest fans run directly from an AC mains source but they are also the least efficient.

Electronically Commutated Fans (EC-DC or ECM fans) take the AC mains source and convert it to a high DC (direct current) voltage, which delivers a more efficient use of the power source, but significantly increases cost. Low voltage DC fans use a power supply to convert the AC mains voltage to a low voltage to operate the fan. These can be more efficient and cost-effective.

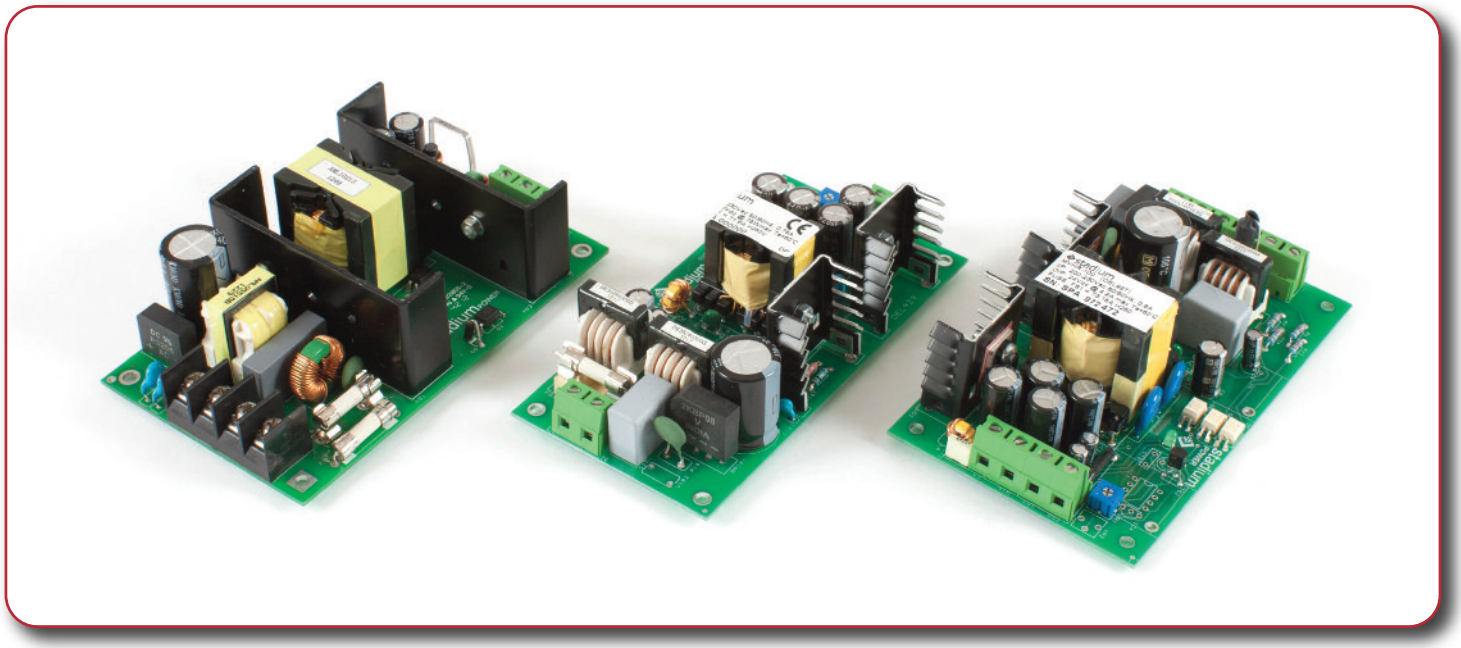

However, the low power and low flow rate requirements of Appendix Q can have a negative impact on the efficiency of Electrically Commuted Fans, which are normally designed to run most efficiently over a given flow rate range.

The problem is actually in how the fan converts the power into mechanical movement; as ever, cost is a factor so the power conversion stage is normally designed to be efficient over a limited window of operation. With MVHR systems, the performance of the fan is largely dependent on the prevailing conditions, number of (wet) rooms and even seasonal climate.

As a result, the most efficient MVHR systems – those that can deliver low power consumption and high efficiency (<1.5W/L/s and >70% heat recovery) over a wide operating window – need a more efficient and typically bespoke power conversion stage. In response to this growing need power supply manufacturers such as Stadium Power are now developing power supplies that meet these exacting requirements.

Given that a typical dwelling may have one kitchen and between 1 and, say, five wet rooms, and each of these rooms may have variable ventilation needs (perhaps controlled by external sensors such as humidity, temperature or occupancy), the load on the power supply could vary between 25% and 100%. For this reason, Stadium Power has designed a range of MVHR PSUs to deliver greater than 90% efficiency at 25, 50, 75 and 100% load, achieving >87% between 10 and 100% loads. Stadium Power now offers five variants in a range developed specifically for MVHR applications; delivering between 75W and 120W @ 24V DC with a 'flat' efficiency of >90% between 25 to 100% load.

MVHR systems are becoming increasingly necessary as we near the Zero Carbon goals for dwellings by 2016. However, looking at system efficiency requires MVHR manufacturers to consider the environmental and financial impact of the variable power requirements such systems exhibit.

In order to not only comply with the requirements of SAP Appendix Q but to also differentiate themselves amongst the competition, MVHR manufacturers should look closely at the PSU sub-components and strive to provide the highest possible efficiency over a wide operating window; parameters many ECM fans aren't designed to deliver.

Through designing-in high efficiency power supplies designed specifically to help meet SAP Appendix Q, MVHR manufacturers can deliver the right solution in the face of mounting demand and competition.

Author profile: Mike Farrer is the Technical Director of Stadium Power

More from Stadium Power
Return to contents page.

Electronic Specifier design

Verifying compliance of USB 2.0 interfaces

As the USB interface remains, by far, the most dominant interface in mobile devices, electronic consumer goods and industrial products, demand for reliable, fast test solutions in development and integration is equally high. By Guido Schulze

In high-speed (HS) mode, the USB 2.0 interface achieves a data transfer rate of 480Mbit/s, and in the PC world (where the interface originated), the 5Gbit/s USB3.0 standard has already become established (super-speed mode). For many mobile applications, PC peripheral devices, equipment and systems used in industry, medicine and A&D environments, the data rate provided by the USB 2.0 interface is fully adequate, which explains its growing popularity in these sectors.

The USB 2.0 standard was published in 2000, and commercially available components based on this standard have now reached technical maturity. Nevertheless, developers are still faced with challenges related to the integration of these components on printed boards and in devices. For example, insufficient isolation with respect to other boards can result in ground loops or crosstalk, which can have a negative impact on the operation of the USB interface. For debugging and stability tests, developers must rely on test solutions that conform to the relevant standards. Accordingly, the organisation responsible for USB standardisation – the USB Implementers Forum (USB-IF) – has defined a test process with appropriate compliance measurements whose purpose is to ensure the error-free interoperability of a wide variety of devices with USB interfaces. Products bearing the USB logo must pass this compliance test, and so developers need automated USB test solutions in order to prepare for compliance testing performed by certified test labs.

Automated compliance

The high measurement accuracy of the R&S RTO oscilloscopes provides a solid foundation for reliable compliance test results. For testing USB 2.0 interfaces in HS mode, either the RTO1024 oscilloscope (2GHz bandwidth) or the RTO1044 (4GHz) can be used. ScopeSuite is a powerful software tool from Rohde & Schwarz that guides the user step-by-step through compliance tests, configures the oscilloscope, automatically executes the measurements and compiles the results in a clear measurement report. The R&S RTO-K21 software option includes USB 2.0 compliance tests for USB devices, hubs and hosts.

To connect the device under test (DUT) to the oscilloscope, Rohde & Schwarz offers the R&S RT-ZF1 test fixture set that can be used for USB 2.0 signal quality tests and for legacy tests on USB 1.1 and USB 1.0 interfaces.

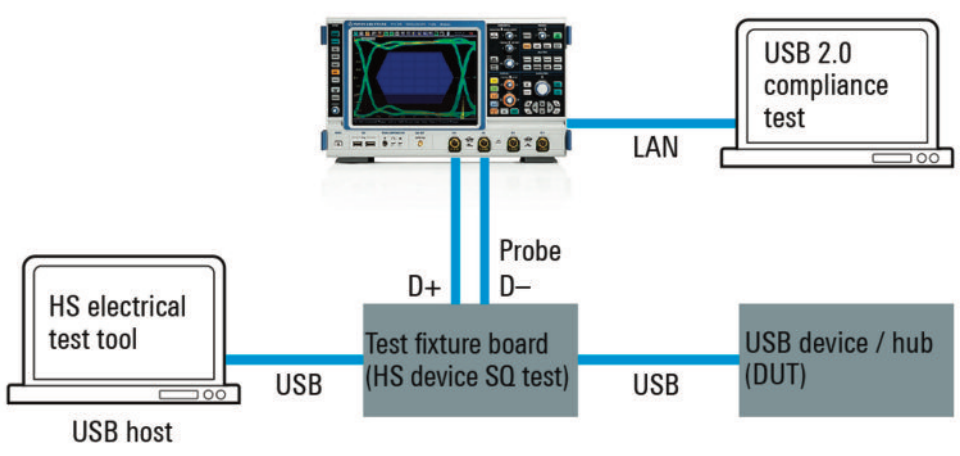

Figure 1: Test setup for HS device signal quality (SQ) test

The test setup varies based on whether the DUT is a USB device, host or hub and on the speed mode selected for compliance testing. The test fixture boards (signal quality board and load board) contain different sections for the individual tests. Figure 1 shows the test setup for HS device signal quality (SQ) tests. R&S ScopeSuite runs on a PC that controls the the oscilloscope, which is connected to the DUT via a differential probe and the test fixture board. The USB-IF software (HS electrical test tool) sets the DUT to the necessary test state. This software should be run on a separate PC since it modifies the USB stack during operation.

ScopeSuite controls the measurement settings and test sequence on the oscilloscope via the LAN interface. Prior to starting the test, the user can define user data, all settings for the test setup and measurement reports. The limit

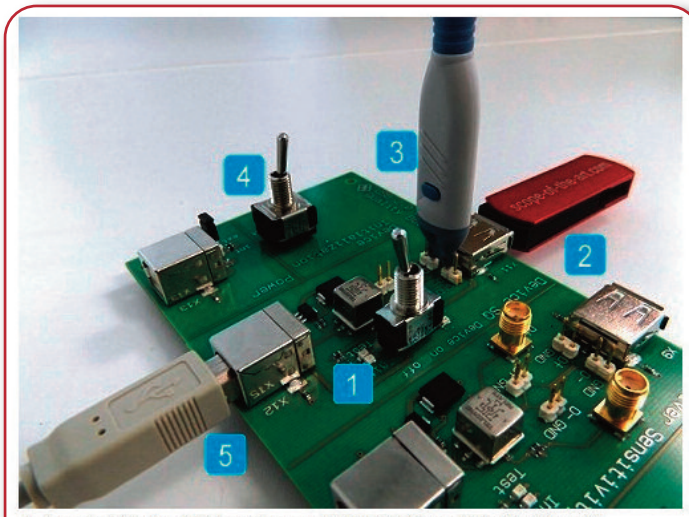

1. Locate "Device SQ" section on "USB 2.0 Signal Quality Board"
2. Insert DUT into X11
3. Connect differential probe from CH 2 to D+/D- nearest to DUT
4. Switch to "Init" position on "Device Initialization" section
5. Attach USB cable between test bed computer and X12

Figure 2: Picture-based instructions for the test sequence: device HS mode, signal quality test with the RT-ZF1 test fixture set

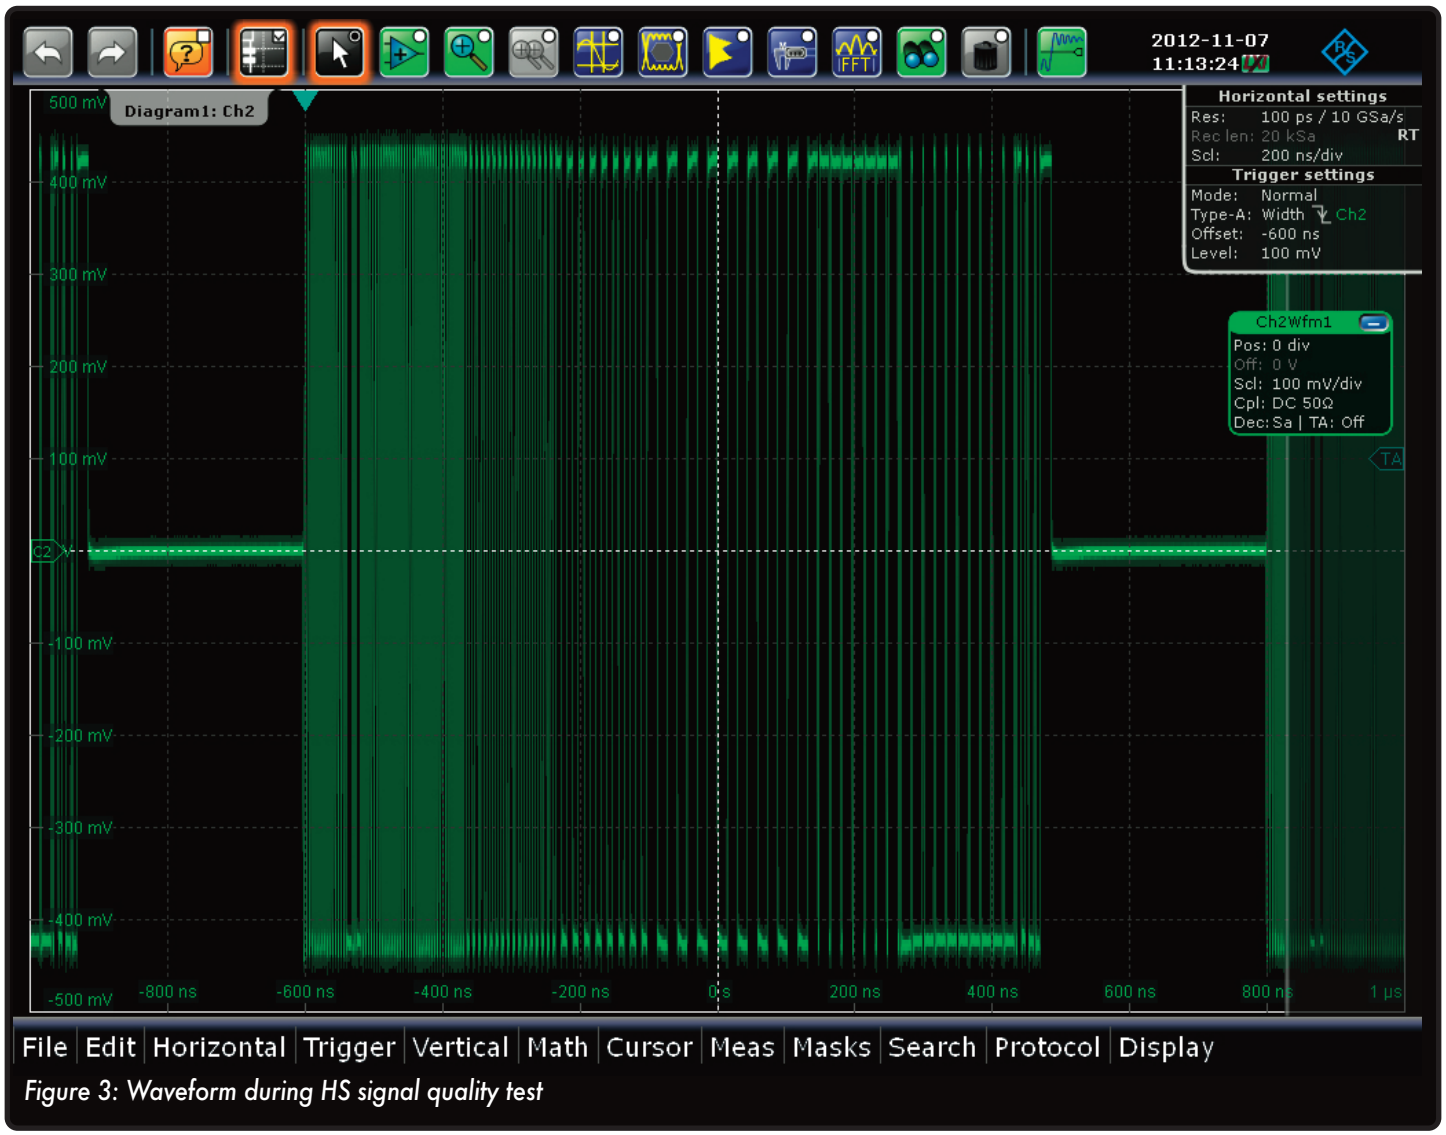

Figure 3: Waveform during HS signal quality test

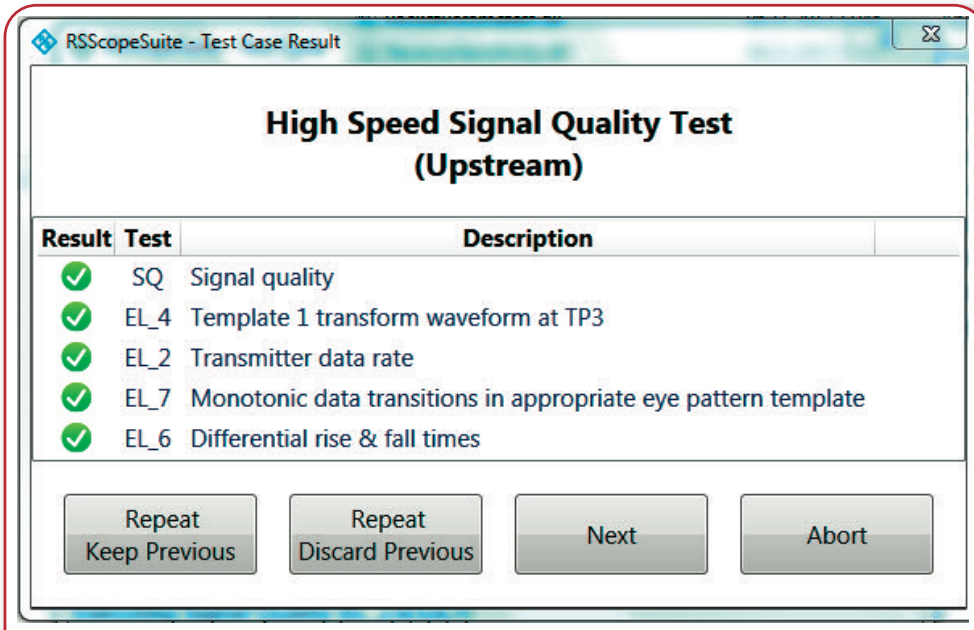

Figure 4: Flexible control of test sequence

editor allows the user to individually adjust the standard-specific test limits. Taking the test setup into account, the software guides the user through all of the selected compliance tests.

High Speed Signal Quality - EL_4

Description	Template 1 transform waveform at TP3
Run	1
Result	Pass
Time	11/07/2012 11:22:16

Additional Information

Measurement	Value	Limits
Signal Eye	Pass	Meet Tpl 1 Tx Wfm Rqmt
Consecutive Jitter RMS	27,2 ps	
Paired JK Jitter RMS	23,971 ps	
Paired KJ Jitter RMS	27,589 ps	

High Speed Signal Quality - EL_4 - Signal Eye

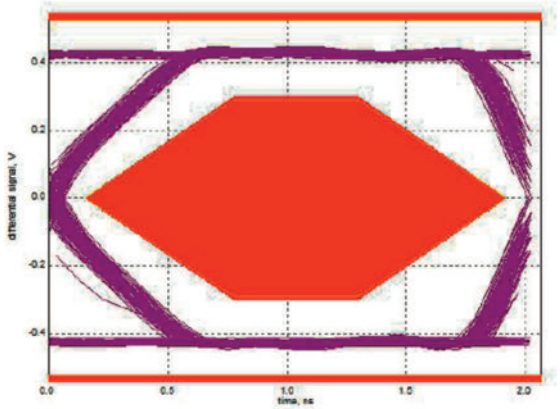

Figure 5: Reports can be individually configured: numerical results and diagram of the eye pattern test

Detailed instructions with pictures make it easy to correctly connect the probes to the test fixture and DUT (Figure 2).

At the beginning of a typical test, the HS electrical test tool configures the test mode for the DUT. The DUT then transmits specific test signals that are captured by the oscilloscope (Figure 3) and passed on to the software for analysis, using the USB-IF electrical test tool to analyse the results.

Test sequences are very flexible with ScopeSuite. For example, the user can use the Repeat – Keep Previous function (Figure 4) to repeat test cases as required for debugging or stability tests. All results are documented in the measurement report. If the user makes an error during the test, such as an improperly connected probe or a test mode that was incorrectly configured with the HS electrical test tool, the result can simply be discarded and the test case repeated using the Repeat – Discard Previous function.

Detailed documentation of measurement results is not only an essential component of compliance testing, it is also very important for debugging and exchanging data with colleagues and customers. ScopeSuite offers an extensive range of documentation functions. The user can, for example, add measurement details and screenshots to the pass/fail results. The software also allows new tests to be added to an interrupted test sequence so that all results will be available in a single report (Figure 5). Reports can also be generated at a later point in time and distributed in various formats including PDF, RTF and HTML.

Author profile: Guido Schulze is Head of Product Management, Oscilloscopes, with Rohde & Schwarz in Munich

More from Rohde & Schwarz
Return to contents page.

Bridging the data deluge gap

Integrating solid state memory in to enterprise system architectures offers benefits in meeting demand for greater storage but comes with challenges. Can flash technology be the answer? By Thomas Pavel

There's no question that RAID is ubiquitous today; the backbone storage technology in modern day datacentres. It's likely to remain so for the foreseeable future with so much cloud computing being delivered by datacentres using RAID technology and by the massive amounts of data being generated by phenomena such as social platforms, smart clients and mobile Internet devices, big data, and the proliferation of video across consumer and enterprise platforms. The adoption of enterprise flash-based solutions is the next logical evolutionary step in storage technology as the gap between data growth and IT infrastructure investment has widened over the years, creating (performance or I/O) bottleneck problems for mission-critical applications.

So how can organisations bridge these gaps? For many, solid state drives are well-placed to close the gap between compute and storage performance because they have a much faster random access time and data transfer rates, lower latency and consistent read performance compared to traditional hard disk drives. However, the cost of replacing the entire storage infrastructure with SSDs is impractical and cost-prohibitive, so a more economic option is to integrate flash technology into enterprise system architectures to deliver significant improvements.

RAID evolution

Over the years, a number of standard RAID schemes have evolved, known as levels. RAID 0 improved performance and added storage but had no fault tolerance while RAID 1 enabled mirrored data to be written identically to two drives. RAID 2 and RAID 3 synchronised disk spindle rotation and stored sequential bits and bytes on a parity drive. With RAID 4, files were distributed between multiple drives that operated independently, allowing I/O requests to be performed in parallel. But with all parity data stored on a single drive, it could suffer from a performance bottleneck. RAID 5 distributed parity along with the data. If a failure occurs, subsequent reads can be calculated from the distributed parity. RAID 6 provides fault tolerance of two drive failures, making larger RAID groups more practical for high availability systems.

Distributing data across multiple drives in a RAID system can be implemented at a software level or a hardware level. At the hardware level, RAID controllers can support many operating systems because they present the RAID as another logical drive. They include a read/write cache, so they can improve performance and as the read/write is non-volatile, pending writes are not lost in the event of power failure, as long as the cache is protected by a backup mechanism. Hardware RAID provides guaranteed performance and does not add computational overhead to the host computer, but because controllers use proprietary data

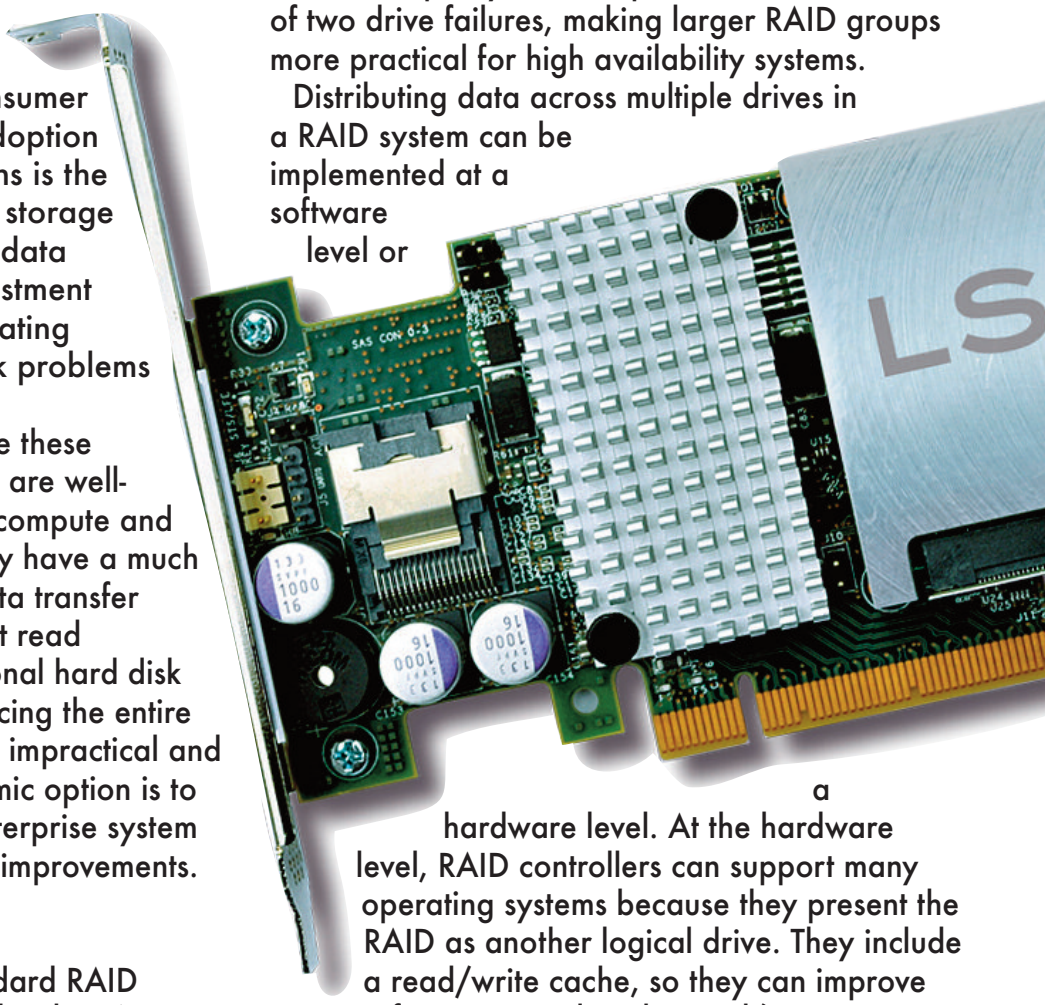

layouts, it may not be possible to work with controllers from multiple suppliers.

Bridging the gap

So far, the evolution of RAID has generally managed to keep pace with the demands placed upon it, but for how much longer? According to research from Cisco, network traffic is likely to experience a compound annual growth rate of 32% between 2011 and 2015. IDC predicts storage capacity will need a compound annual growth rate of 50% in the same period. These trends are emerging at the same time as Gartner has estimated the compound annual growth rate for IT spending is 5% and the figure for telecom equipment spending is 7%.

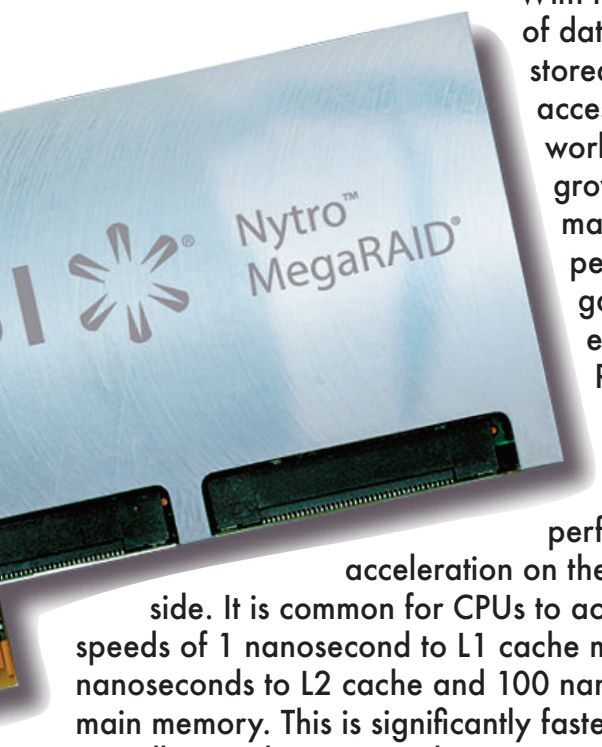

With the amount of data being stored and accessed worldwide growing at a massive rate, a performance gap is emerging as RAID struggles to keep up with performance acceleration on the processing side. It is common for CPUs to achieve write speeds of 1 nanosecond to L1 cache memory, 10 nanoseconds to L2 cache and 100 nanoseconds to main memory. This is significantly faster than the 10 millisecond write speed to tier 1 storage, or the 20 milliseconds to tier 2 storage and to near line storage. This represents a latency penalty of 100,000x for leaving the memory hierarchy.

The combination of this performance gap and the explosion in data growth and network traffic is likely to strain RAID storage infrastructure, creating bottlenecks, throttling application performance and making it harder for companies to extract the full value from their data. This is of even greater concern in a world where speed of access to data is of the essence and everybody wants access to everything immediately.

Redefining performance

Using flash-based storage with existing storage can save significant amounts of money because tiered storage arrays with flash on the RAID controller (combined with intelligent software) can replace a huge amount of disk drives that are otherwise necessary to maintain I/O rates in traditional storage arrays. The combined approach enables businesses to intelligently use flash storage and their existing hard drives together in a way that can give them optimum price/performance in a tiered storage environment.

Combining PCIe flash technology with intelligent caching and management software can deliver an impressive performance acceleration that depends on configurations and application but is easily a large multiple of the HDD performance. It is not unusual to experience application performance acceleration of 5-6x but some cases have been reported up to 30x.

Solutions can be delivered with different capacities to suit different requirements. High capacity PCIe flash solutions can be used for primary storage to deliver high value non-transparent storage, but have a cost premium attached. Medium capacity flash solutions that combine PCIe flash technology with intelligent caching software can accelerate high value SAN and complex DAS connected storage and provide a balance between cost and value. Low capacity flash solutions combine a RAID controller card with on-board flash and intelligent caching software to accelerate DAS connected storage and bring the value of flash to the masses.

The intelligent deployment of flash technology can help RAID evolve to meet the challenges presented by the 'data deluge gap'. And perhaps it's fitting, given the preoccupation with performance that informed the publication of 'The Case for Redundant Arrays of Inexpensive Disks' in 1988, that flash technology is also being deployed to help RAID address I/O performance issues.

Author profile: Thomas Pavel is the Director of Channel Sales at LSI

More from LSI Europe
Return to contents page.

Hybrid drives capacity and performance

The hybrid drive, which combines the cost advantage of rotating magnetic disc storage with the high performance of NAND flash memory, has recently been attracting attention from computer users looking for fast, large capacity drives.

By Mine Budiman, Eric Dunn & Rick Ehrlich

The combination of a hard disk drive (HDD) and NAND flash memory has the potential to deliver a solid-state drive (SSD) like user experience. Hardware, firmware and architecture design are all essential to creating a product that will meet user expectations in terms of performance, capacity and cost. In the case of the latest technology, for example, that includes developing a new cache algorithm that places frequently used data into NAND flash memory, resulting in a hybrid drive with SSD-like performance.

Large numbers of video and other types of data are being handled more frequently than ever before. People want to store more data on devices installed in PCs (e.g., HDDs), yet still achieve increased data access performance. For example in operating system (OS) data readout performance for quick booting PCs, as well as improved data read/write performance while PCs are in operation. HDDs are still the major player in the storage device arena and are used in numerous fields, though there are limitations on performance improvement potential. For this reason, SSDs, which feature high-speed data access, are attracting attention and the SSD market is expanding. However, HDDs still have the advantage in cost per gigabyte. Given these circumstances, Hybrid drives, which combine HDDs with NAND memory, have come into their own as storage devices capable of realising both high capacities and high performance. Now, by adopting an intelligent cache algorithm that learns what data should be stored in the NAND memory cache, the latest generation of these drives is set to deliver HDD capacity and SSD-level performance to both enterprises and consumers.

Figure 1: The number of hybrid drives built-in to notebook PCs is predicted to increase greatly by 2015

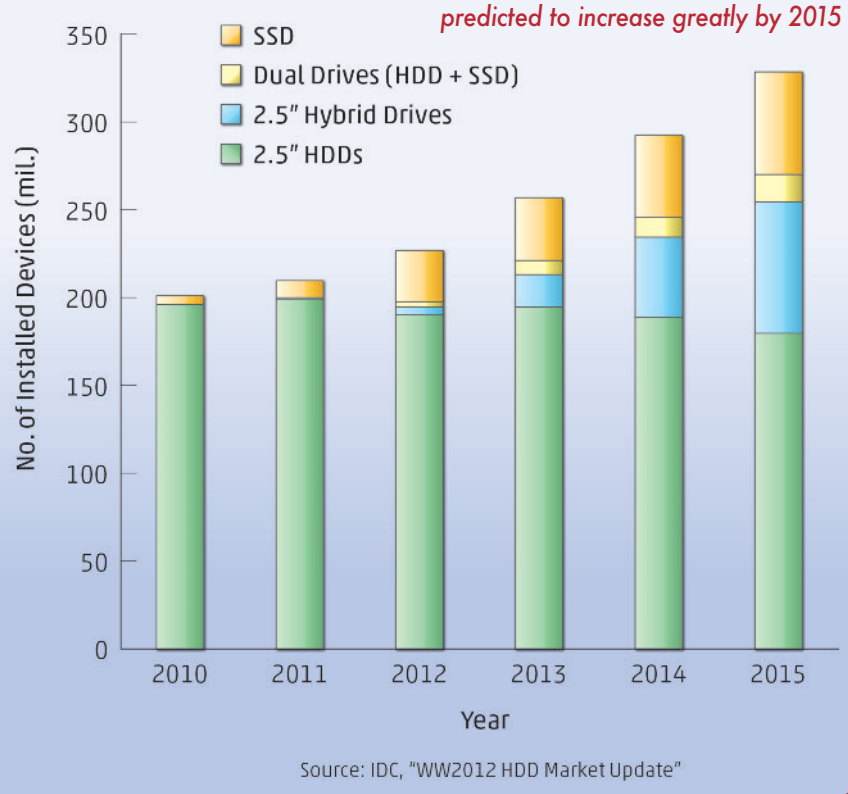

The hybrid drive achieves both a high capacity and high performance by combining traditional NAND memory technologies, NAND memory handling technologies (instrumental in SSD technologies), and HDD product development technologies. Full-fledged adoption of hybrid devices are estimated to account for about 25% of the market in 2015 (Figure 1).

In the basic structure of a hybrid drive, the cache memory, the HDD uses low-capacity, high- speed

DRAM (volatile memory) and a large-capacity magnetic disk (non-volatile memory). NAND memory (non-volatile memory), which has medium capacity, read/write speeds and bit costs compared to DRAM and magnetic disks, is added as a secondary cache in this HDD.

A newly developed cache algorithm dynamically studies data access patterns and stores frequently accessed data in non-volatile NAND memory. This allows access performance to be improved even if no cache hits occur in DRAM during use, for example when booting the PC. Thus, the algorithm can help the hybrid drive achieve performance levels close to that of SSDs.

PC systems recognise the hybrid drive just as they do conventional HDDs and SSDs. There is therefore, no need for additional software and existing HDDs or SSDs can be replaced directly with this drive.

Hardware development

In order to keep hardware costs low while achieving high performance, Toshiba has developed a NAND controller for managing data read/write to and from NAND memory. This controller serves as a bridge chip positioned between the System on a Chip (SoC) and NAND chip. The NAND controller and Synchronous DRAM (SDRAM) are connected in parallel to the Double Data Rate 2 (DDR2) SoC interface; different address spaces are allocated to the NAND controller and SDRAM to allow the existing HDD and NAND chips to share the SoC.

To avoid totally redesigning the HDD the aforementioned NAND controller and NAND memory need to be mounted onto the HDD circuit board. The conventional four-layer printed circuit board has been changed to a six-layer board allowing the components to be mounted on the same size board as that used by the standard HDD as shown in Figure 2.

The amount of space needed to mount the components was further reduced by adopting small sized DDR2 (rather than DDR1) for SDRAM. However, it was impossible to mount more than one NAND memory module. The candidate NAND memory modules and a performance comparison of sequential write operations are shown in Table 1. With respect to NAND memory, single level cell (SLC) technology has a higher cost

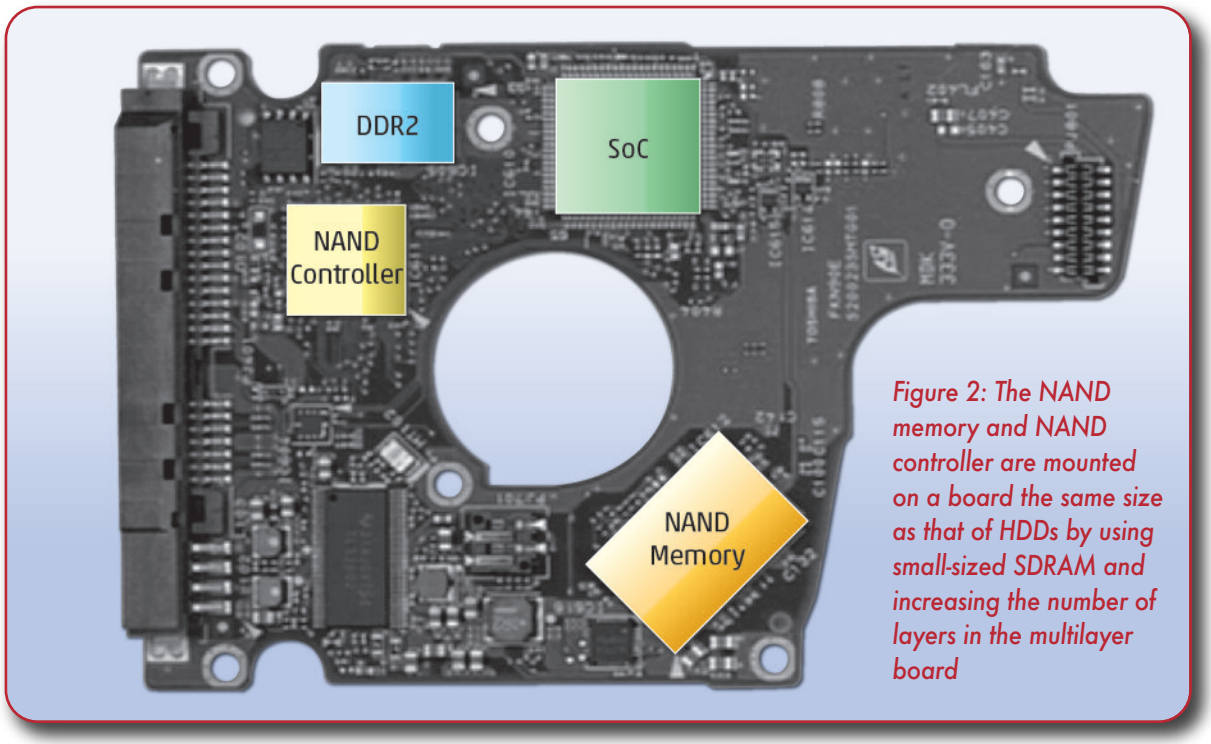

Figure 2: The NAND memory and NAND controller are mounted on a board the same size as that of HDDs by using small-sized SDRAM and increasing the number of layers in the multilayer board

of capacity per unit than multi level cell (MLC) technology, but also offers much higher performance. If the number of dies (NAND chips) is increased, the drive's performance will be improved considerably because of parallel processing. Since this hybrid drive could not include more than one NAND memory, a cost/performance trade-off analysis was performed, which selected the 8GB SLC (which uses four 2-GB dies).

In typical HDDs, random access performance is restricted by seek-times and rotational delays; in most cases, command reordering techniques are used to minimise the impact of these delays, but HDDs are still limited to 200-300 Input/Output per second (IOPS). NAND memory performance is restricted by the access speed to the NAND memory, but this can be improved by increasing the number of channels and the number of dies to enable parallel access. With such techniques, it is possible to enhance the random read/write performance of NAND memory up to several thousand IOPS, an order of magnitude better than that of a standard HDD.

Software development

The capacity of the NAND memory mounted on the hybrid drive is 8GB, approximately 1% of the capacity of a 1TB HDD. Effective use of NAND memory as a cache requires careful consideration of what data to place in the NAND memory because of the limitations imposed by its capacity. The hybrid drive's design concept is shown in Figure 3.

The hybrid drive consists of the following three storage tiers: SDRAM, NAND memory, and magnetic disk, each of which has different read/write speeds. When the system requests that data be written, the hybrid drive will write the data to the NAND memory before writing the data to the magnetic disk (data will be written to disk later). When the system requests that data be read from a specific address, the drive will check whether the data for the specified address is stored in NAND memory. If the data is present in NAND memory, the drive will quickly send the data to the system. If not, the drive will read the data from the magnetic disk and send the data to the system, then copy the data to NAND memory. Because of this algorithm, data that was recently written to the Hybrid drive or read out from the drive is contained in NAND memory. If the system requests that data be read from the same address again, the drive can return such data to the system very quickly because the required data is stored in NAND memory.

Using a hybrid drive prototype with a 16GB NAND memory a series of six iterations was conducted with the performance-benchmarking tool (PCMark Vantage HDD Suite). As for the benchmark results, the relationships between benchmark scores and magnetic disk read counts are shown in Figure 4. The Hybrid drive must read data from the magnetic disk if the data is not stored in NAND memory. After the second measurement, the magnetic disk read count decreases dramatically, since most of the data is

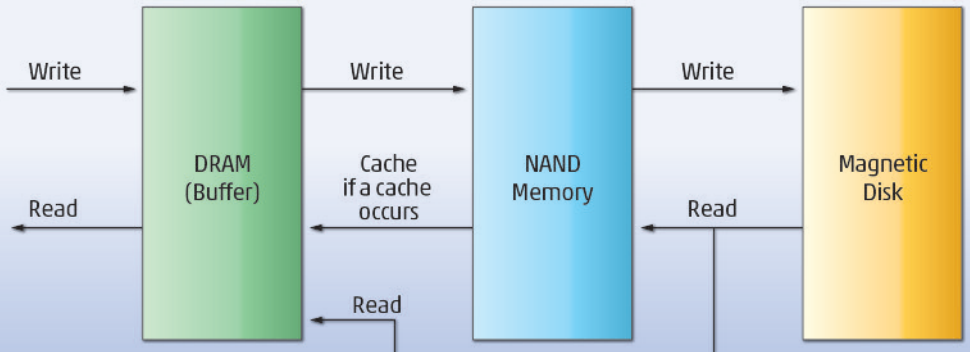

Figure 3: Data is controlled so that recently written and read out data is stored in NAND memory, allowing read/write speeds to be improved

read out from NAND memory. In the first measurement, the benchmark score is only slightly better than that of a conventional HDD, but the scores for subsequent benchmarks exceed those for the first measurement by a factor of three. These scores are the result of studying the system's data access patterns and caching the necessary data into NAND memory.

In an example time comparison, in which an HDD, Hybrid drive, and SSD were installed in three otherwise identical PCs, the time to load the same application was measured. The load time required for the Hybrid device is much less than that of the conventional HDD and is roughly equal to that of the SSD.

The development of a hybrid drive containing three storage levels comprising SDRAM, NAND memory, and magnetic disk answers the need for many users of fast, high capacity storage. By utilising an algorithm that controls how and where data is stored and cached, SSD-comparable performance can be achieved for repetitive access cycles such as OS access times during PC boot sequences.

To further improve hybrid drive performance, a specification for standardisation of hints, which are data provided by PCs that give hints as to whether to keep data in NAND memory, is now being discussed at the Serial ATA International Organization (SATA-IO), the standard-setting organisation for Serial Advanced Technology Attachment (SATA). In the NAND memory field, efforts to reduce costs by integrating components and miniaturising processes are currently underway.

Toshiba continues to improve the cache algorithm to intelligently exploit the 3-tiered structure of SDRAM, NAND memory, and a magnetic disk, and also to develop second-generation hybrid drives with even higher performance.

Figure 4: Benchmarks show a score increase of about 3x of a conventional HDD

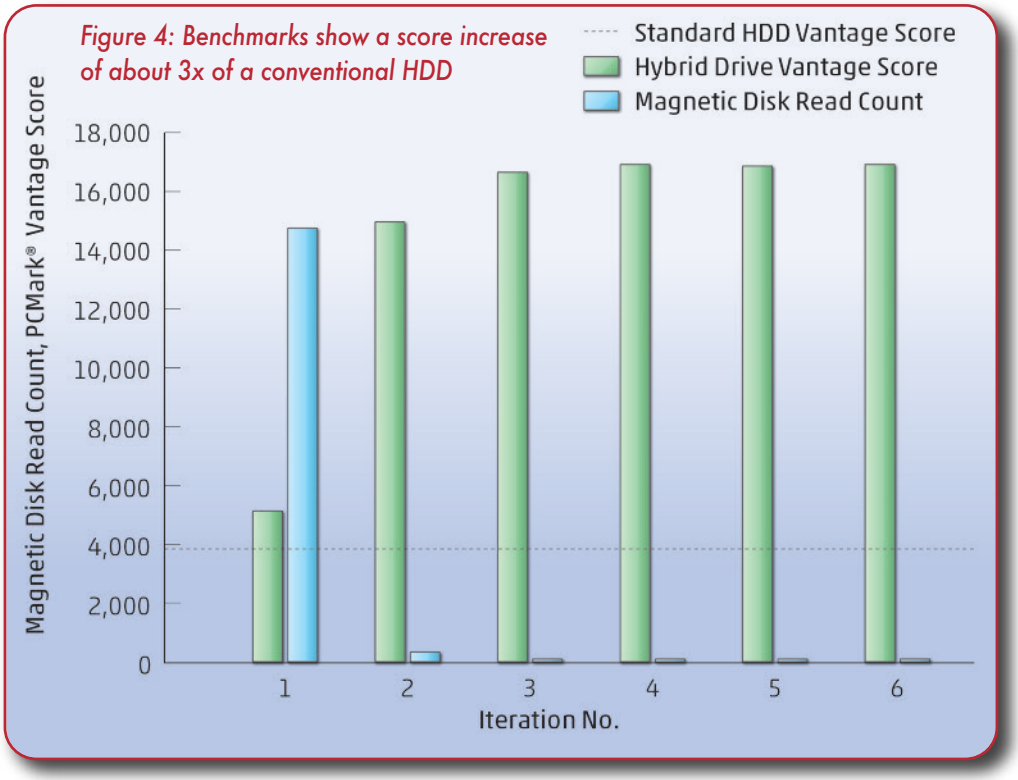

Author profiles: Mine Budiman, Eric Dunn & Rick Ehrlich work at Toshiba's Storage Products Division

More from Toshiba Electronics Europe
Return to contents page.

Everyone's an individual

How a new approach to the single-strand wiring of PCBs addresses existing challenges by offering a powerful, reliable, economical and easily installable interconnection solution. By Kevin Canham

The trend towards the miniaturisation of the electronic circuitry used in industrial systems means that more functions are being combined in single devices or applications. At the same time, individual devices are becoming more flexible to fulfil different functions such as memory, measurement and control, both centrally and peripherally. Powerful miniature devices require powerful termination concepts that can also be implemented economically.

In the future, industrial systems, machines and instrumentation will continue the trend towards further miniaturisation. As a result, more and more functions will have to be combined in one device or one application. Application software will take on additional functions and will become increasingly powerful. At the same time, individual devices will become smaller and smaller, with no significant reductions in performance and reliability.

A not insignificant reason for these trends involves the implementation of electronic control elements that make the devices more flexible and their use more precise. Devices that started off as specialised or 'insular' systems are now becoming suitable for broader use. At the same time, they can be implemented in complex systems where they fulfil different functions as memory, measurement or control devices, both centrally and peripherally. Devices and systems that can be used with increasing precision are able to react to environmental conditions and can be controlled in the context of the end application.

The advantages that industry, users and consumers expect from these developments speak for themselves: the massively increased power densities and more compact dimensions of the devices result in broader application areas and more possible uses. Today, infinitely adjustable controls combined with powerful data-acquisition functions are expected: even for simple applications that until recently involved only a simple 'on/off' circuit.

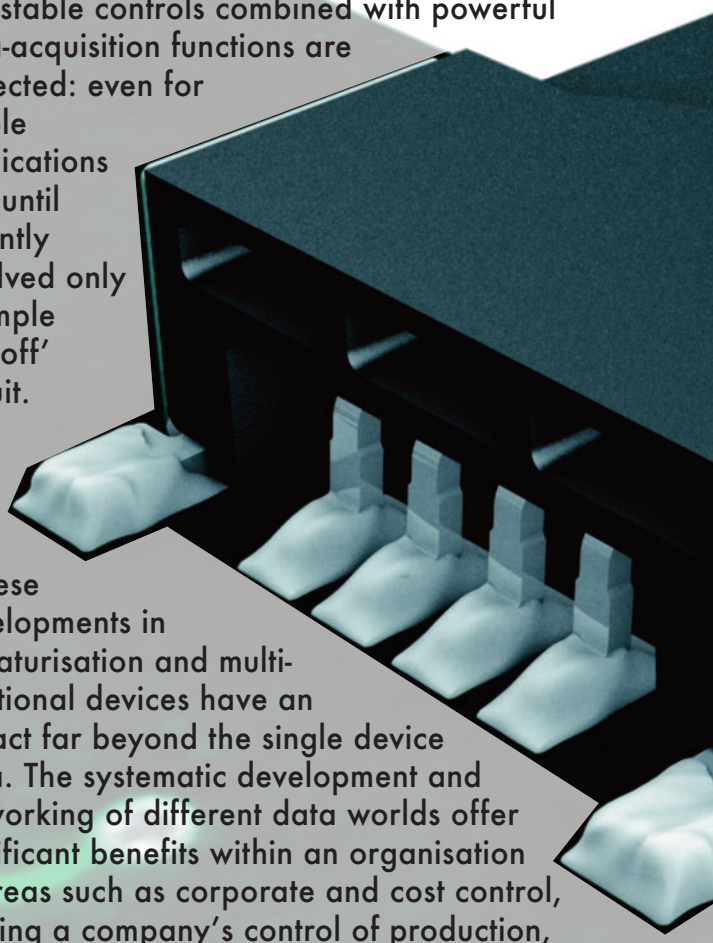

These developments in miniaturisation and multi-functional devices have an impact far beyond the single device area. The systematic development and networking of different data worlds offer significant benefits within an organisation in areas such as corporate and cost control, helping a company's control of production, energy and cash flows. Service and maintenance work can be arranged more effectively, and the associated costs can be minimised.

Modular concepts

One result of this trend is an increasing use of modular devices, which in turn present challenges in terms of field installation, service and maintenance. Even in conventional service and maintenance work, the increased use of modules and devices in the field that must be matched by the capability for swift and simple

field installation and exchange without special operator training or the need for special tools.

These challenges extend down to the level of individual devices and into peripheral areas such as the printed-circuit board and the termination technology required to link the device to the PCB. Powerful miniature devices require powerful termination concepts which, most importantly, need to be implemented economically. Customised production, complex manufacturing techniques or complicated assembly methods that may be subject to human errors in the field cannot survive in an economy that is based on the effective use of resources, and that favours economical solutions.

In order to address these challenges, The Harting Technology Group has investigated new concepts and solutions for single- strand wiring of PCBs that are also powerful and reliable, economical and easily installable in the miniature range. The result is an interconnection concept known as har-flexicon.

Using this device, it is possible to connect flexible single conductors using insulation displacement connection (IDC) technology without the need for separate wire-stripping. This process allows simple disconnection or connection of the connector without the need for complex re-contacting. With a pitch of 1.27mm, the connector's height is only 5mm, measured from the upper edge of the PCB. In spite of its compact design, the har-flexicon can be operated very simply and swiftly. The mating reliability is extremely high thanks to the use of error-free contact retention, and multiple insertion and detachment steps are very straightforward.

The use of IDC connection technology makes field installation possible without the need for special tools

No tools are required to detach or insert the har-flexicon in the field. This reduces costs significantly, especially in the maintenance and service area, and as the connection can also be executed swiftly under difficult environmental conditions.

This is currently the most compact field-installable PCB connector available for use with single wires. The 1.27mm pitch is only half that used with conventional termination technology and the device is also suitable for universal use in transferring I/O signals in sensor systems and actuating elements.

User benefits

This concept for the connection of PCBs in industrial devices offers several advantages over conventional solutions, including swift single-strand termination for flexible processing in the field without the need for special tools, compatibility with surface-mount technology, miniaturised pitch size and a high stability level.

The har-flexicon is a logical continuation of Harting's har-flex product family, which provides rugged and economical board-to-board connection technology. As a new generation of PCB connectors, both series are setting the future standard for the efficient PCB connection and rugged miniaturisation that state-of-the-art devices require in industrial applications. Together, these connector series provide a complete range from device connectivity to device internal wiring to board-to-board connections. A standard connector design results in uniform devices.

The new termination concept can be processed by machine and can be applied to the PCB using the same surface-mount technology used for all other components on the PCB without any compromise on stability. In other words, with the PCB now becoming the core component in industrial devices, standard machine processing of the components will become the accepted norm.

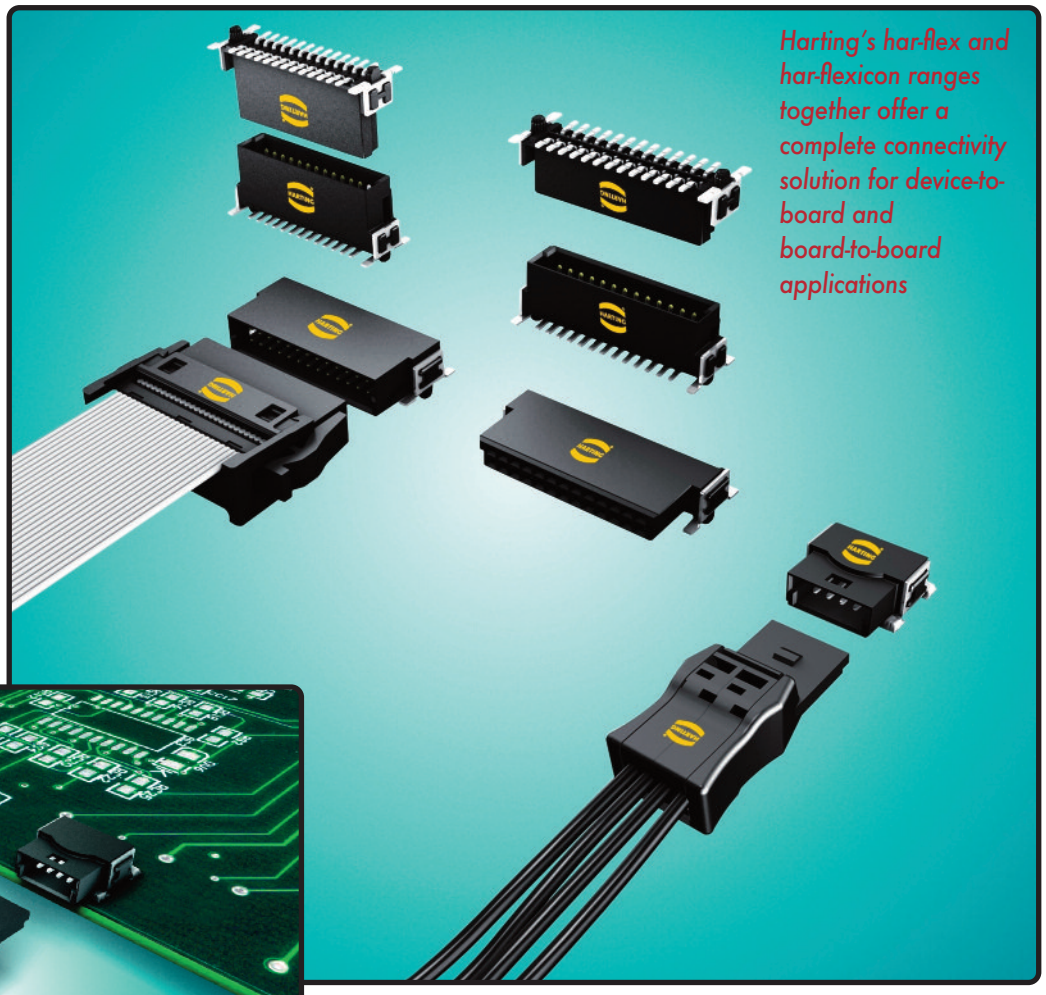

Harting's har-flex and har-flexicon ranges together offer a complete connectivity solution for device-to-board and board-to-board applications

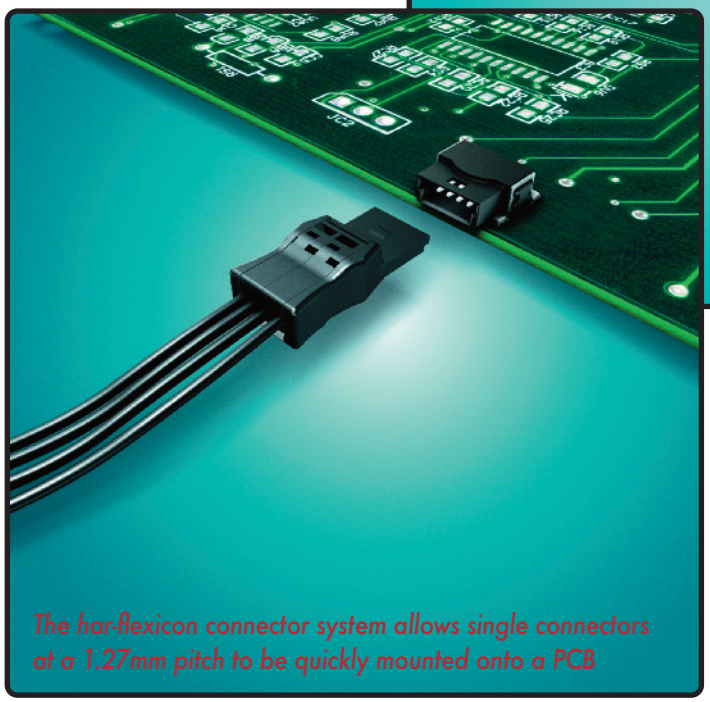

The har-flexicon connector system allows single connectors at a 1.27mm pitch to be quickly mounted onto a PCB

Traditionally, electronic components such as passive devices or ICs have been mounted on the PCB using surface-mount processing as the standard. Connectors, however, have been the last devices to be mounted using conventional wiring, largely because soldered connections display greater strength for through-hole plating.

To overcome the problem of SMT connections being less stable, the new rugged connector concept uses extensive lateral surface-mount fixings which allow it to absorb larger pulling and insertion forces. This is a key advantage, particularly for detachable connections in which repeated and simple plugging must be ensured. This approach makes it possible to assemble the PCB terminals in the same operational process and with the same equipment as the other surface-mount devices because the required stability is ensured.

This approach makes it possible to ensure clear advantages through swift and standardised machine processing. It also eliminates the possibility of weak points inherent in the soldering of individual wires, such as limitations of the heat input on to the PCB and fluctuations in the processing quality.

It also becomes possible to use serial machine processing with the SMT components being delivered in tape packaging. The high temperature materials in the insulating inserts also allow safe processing in the reflow process.

Author profile: Kevin Canham is the Production and Application Manager with Harting Technology Group

More from Harting
Return to contents page.

Board Mounted Power

- ✓ Low cost
- ✓ High reliability
- ✓ Short lead times
- ✓ Industry standard footprints
- ✓ Highest power density

AC-DC low power board mounted ranges

ECE series
- 5 to 40 Watts
- Ultra compact from 1” x 1”
- Encapsulated
- Single and dual outputs
- <0.3 Watts no load input power

ECL series
- 5 to 30 Watts
- Ultra compact size
- Single & multi outputs
- Open-frame & encapsulated
- <0.3 Watts no load input power

Industry standard DC-DC ranges

1 to 3 Watts
- SIP & DIP packages
- Regulated & unregulated versions
- ±10%, 2:1 & 4:1 inputs
- 500 V to 6 kV isolation

2 to 60 Watts
- 1 x 1”, 1 x 2”, 2 x 2” & DIP-24 footprints
- Regulated single, dual & triple outputs
- Wide 2:1 & 4:1 inputs
- Up to 3.5 kV isolation

Medical safety approvals
- 3 to 10 Watts
- Meets IEC60601-1 3rd edition
- 4000 VAC reinforced insulation
- 2µA patient leakage current

LED drivers
- 5 to 48 Watts
- Up to 1000 mA output
- Constant current output
- 95% efficiency

Surface mount power
- 1 & 2 Watts
- Regulated & unregulated versions
- ±10, 2:1 & 4:1 inputs
- Single & dual outputs

Download FREE
interactive guides to help select your AC-DC & DC-DC products at:
www.xppower.com/literature_downloads

Farnell Also available from www.farnell.com

Visit our website to request a copy of our latest Power Supply Guide and see our complete line of power products.

GREEN•POWER

Selector App Available

XP Power
www.xppower.com

Shock resistance

Tackling EMC in medical power supplies requires innovative design

All modern power supplies are of the switched-mode type, as these are smaller and more efficient than the old linear types. Switched-mode supplies, however, generate electromagnetic interference (EMI), both conducted and radiated, and require the incorporation of EMI filters.

The leakage current requirements of the IEC60601-1 are difficult to achieve while, at the same time, meeting EMC compliance requirements. The maximum permissible normal condition earth leakage is 300µA for worldwide approvals, but this figure applies to the end-product as a whole, not just the power supply. To allow for additional leakage in other components it is highly desirable for the power supply to have an even lower leakage current.

The capacitors in the EMI filters allow a small amount of leakage current to flow and the more effective the filter at suppressing the interference, the more leakage it is likely to produce. It seems, therefore, that there is a trade-off between EMC performance and leakage current.

This leads to an interesting challenge since EMC performance is another crucial issue for medical power supplies. For conventionally designed switched-mode supplies this is indeed true, but EMC performance can be improved by methods other than simply providing more filtering. A better approach is to minimise the amount of interference that the power supply generates in the first place.

There are also the expanding RoHS (controls hazardous substances) and WEEE (waste controls and recycling) Directives to consider. Up to now 'medical devices' and 'monitoring and control instruments' have been exempt from the RoHS and WEEE Directives. Although these exceptions, for many medical devices and monitoring/control instruments, will be removed, many medical equipment manufacturers have already completed or are in the process of modifying their products to comply with these directives.

Improving medical power supplies

The FETs (field effect transistors) used as electronic switches in modern switched-mode power supplies are usually configured to switch as quickly as possible because this helps to minimise losses. Unfortunately the faster the FETs switch, the more interference the switching circuit generates.

Some of the best modern power supply designs, therefore, deliberately slow down the switching operation using special 'zero-voltage switching' or 'ZVS' circuits so that the power supply's efficiency is not compromised. ZVS circuitry still allows relatively fast switching of the transistors while achieving voltage transitions (rise and fall times) that are much slower – in of the order of 100ns (nano seconds) compared to 20ns in conventional hard-switching power supplies.

In turn, the amount of electromagnetic interference generated is greatly reduced and therefore only a small EMI filter is needed for these supplies to meet the EMC requirements of even the most demanding medical applications. With only a modest amount of filtering needed, leakage currents can also be kept to a minimum, satisfying another important requirement.

A further benefit is that the ZVS circuitry eliminates the need for an inter-winding shield within the transformer, another technique which was traditionally employed to improve EMC performance. Eliminating this shield not only allows a physically smaller transformer to be used, thereby reducing the overall size of the power supply, but also further increases the efficiency.

The majority of switched-mode power supplies designed over the last 10-20 years use hard switching pulse-width modulated circuits; the latest generations however are using resonant and multi-resonant circuits to achieve the highest possible efficiencies. Some manufacturers also offer power supplies with a number of leakage current options to allow OEMs to find the best trade-off between EMI and earth leakage current for their application.

Digital control

The most recent advancement in medical power supply design is the implementation of digital control technologies. A number of manufacturers are replacing analogue 'housekeeping' circuits (under-voltage lockout, fan speed control, customer signals, etc.) with microcontroller based solutions to achieve a reduction in parts count and circuit complexity. Some are also introducing new products that incorporate full digital control of the power supply, which enables improved characteristics, such as a significant peak power rating under all input voltage conditions, further parts count reduction and greater reliability.

The use of digital control (microcontroller-based) allows these power supplies to be smaller in size and much more efficient, consistent with the trend towards environmentally friendly products. So, in addition to medical equipment, these digitally controlled power supplies can be incorporated in industrial and commercial designs where space is limited, providing a smaller and cooler operating end-product.

TDK-Lambda's EFE range of digital power supplies include active PFC (power factor correction) which ensures EN61000-3-2 compliance and operation from a wide-range input from 90 to 264Vac for global applications. Normal condition earth leakage current is less than 300µA at up to 264Vac input, fully complying with most medical safety requirements. Other EMC improving design features include the use of low-loss Silicon Carbide (SiC) Schottky diodes in the PFC circuit which combined with a ZVS topology for the DC:DC conversion ensure Curve B EMC performance with a significant margin. These new digital power supplies are of course certified per the IEC, EN, UL, CSA and ANSI/AAMI 60601-1 standards for medical equipment, and are also be approved to IEC/EN/UL/CSA 60950-1, and designed to meet IEC/EN/UL/CSA 61010-1. The EFE300M digitally controlled power supply meets the rigorous international safety standards for medical equipment, making it suitable for use in B and BF type medical applications.

More from TDK Lambda
Return to contents page.

Opening doors

Introducing the OpenCL standard and how to implement it on an FPGA.

By Deshanand Singh

As multicore processing devices first came on to the market it was recognised that there needed to be a standard model for creating programs that will execute across multiple cores and potentially different devices. The lack of a standard that is portable across different programmable technologies had plagued programmers. It was with some relief that in late 2008 Apple submitted a proposal for an OpenCL (Open Computing Language) draft specification to The Khronos Group in an effort to create a cross-platform parallel programming standard. The Khronos Group consists of a consortium of industry members such as Apple, IBM, Intel, AMD, NVIDIA, Altera, and many others. This group has been responsible for defining all the OpenCL specifications, the most current being Version 1.2.

The OpenCL standard allows for the implementation of parallel algorithms that can be ported from platform to platform with minimal recoding. The language is based on the C programming language and contains extensions that allow for the provision of parallelism.

In addition to providing a portable model, the OpenCL standard inherently offers the ability to describe parallel algorithms to be implemented on FPGAs, at a much higher level of abstraction than hardware description languages (HDLs) such as VHDL or Verilog. Although many high-level synthesis tools exist for gaining this higher level of abstraction, they have all suffered from the same fundamental problem. These tools would attempt to take in a sequential C program and produce a parallel HDL implementation. The difficulty was not so much in the creation of an HDL implementation, but rather in the extraction of thread-level parallelism that would allow the FPGA implementation to achieve high performance. With FPGAs being on the furthest extreme of the parallel spectrum, any failure to extract maximum parallelism is more crippling than on other devices. The OpenCL standard solves many of these problems by allowing the programmer to explicitly specify and control parallelism. The OpenCL standard better matches the highly parallel nature of FPGAs than do sequential programs described in C alone.

Figure 1: Overview of OpenCL

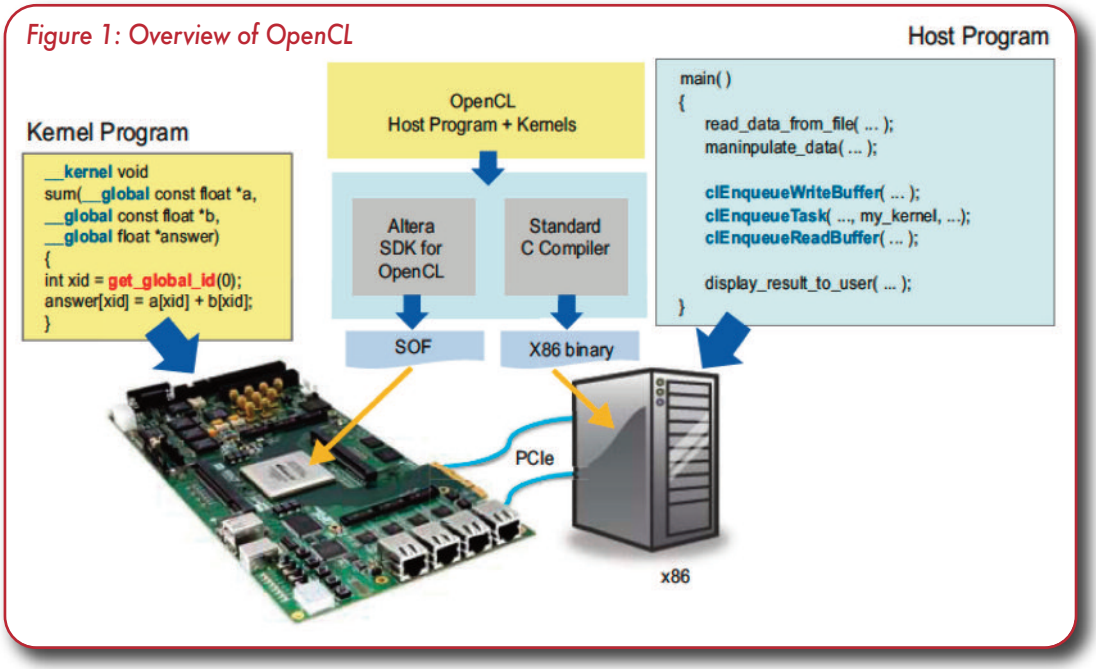

OpenCL applications consist of two parts. The OpenCL host program is a software routine written in standard C/C++ that runs on any sort of microprocessor. That processor may be, for example, an embedded soft

processor in an FPGA, a hard ARM processor, or an external x86 processor, as depicted in Figure 1.

At a certain point during the execution of this host software routine, there is likely to be a function that is computationally expensive and can benefit from the highly parallel acceleration on a more parallel device: a CPU, GPU, FPGA, etc. This function to be accelerated is referred to as an OpenCL kernel. These kernels are written in standard C, however they are annotated with constructs to specify parallelism and memory hierarchy. The example shown in Figure 2 performs the vector addition of two arrays, A and B, while writing the results back to an output array answer. Parallel threads operate on each element of the vector, allowing the result to be computed much more quickly when it is accelerated by a device that offers massive amounts of fine-grained parallelism, such as an FPGA. The host program has access to standard OpenCL application programming interfaces (APIs) that allow data to be transferred to the FPGA, invoking the kernel on the FPGA and transferring the resulting data back.

```
__kernel void
sum(__global const float *a,
__global const float *b,
__global float *answer)
{
int xid = get_global_id(0);
answer[xid] = a[xid] + b[xid];
}
```

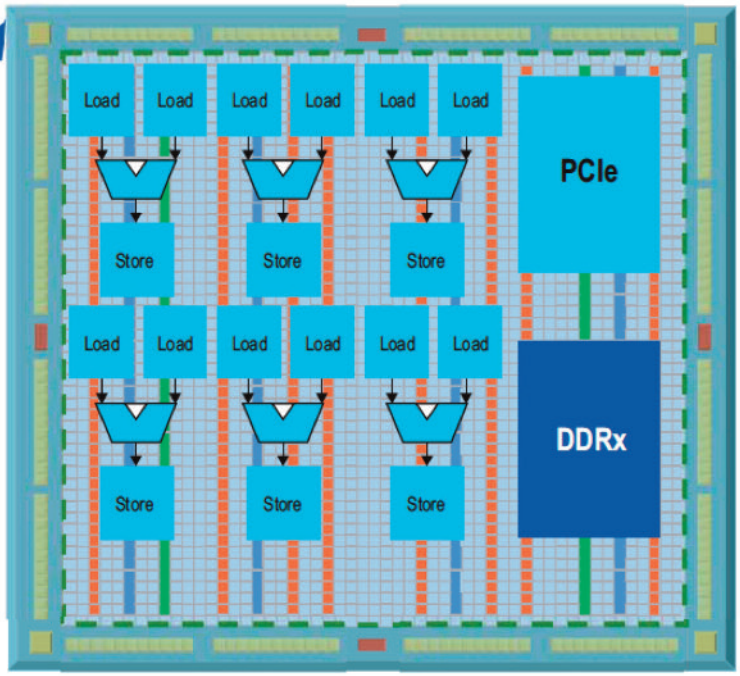

Figure 2: Example of OpenCL on an FPGA

Unlike CPUs and GPUs, where parallel threads can be executed on different cores, FPGAs offer a different strategy. Kernel functions can be transformed into dedicated and deeply pipelined hardware circuits that are inherently multithreaded using the concept of pipeline parallelism. Each of these pipelines can be replicated many times to provide even more parallelism than is possible with a single pipeline. For example, Altera's OpenCL Compiler translates an OpenCL kernel to hardware by creating a circuit that implements each operation. These circuits are wired together to mimic the flow of data in the kernel. In our vector addition example, the translation to hardware will result in a simple feed-forward pipeline. The loads from arrays A and B are converted into load units, which are small circuits responsible for issuing addresses to external memory and processing the returned data. The two returned values are fed directly into an adder unit responsible for calculating the floating-point addition of these two values. Finally, the result of the adder is wired directly to a store unit that writes the sum back to external memory.

The most important concept behind the OpenCL-to-FPGA compiler is the notion of pipeline parallelism. For simplicity, assume the compiler has created three pipeline stages for the kernel, as shown in Figure 3. On the first clock cycle, thread 0 is clocked into the two load units. This indicates

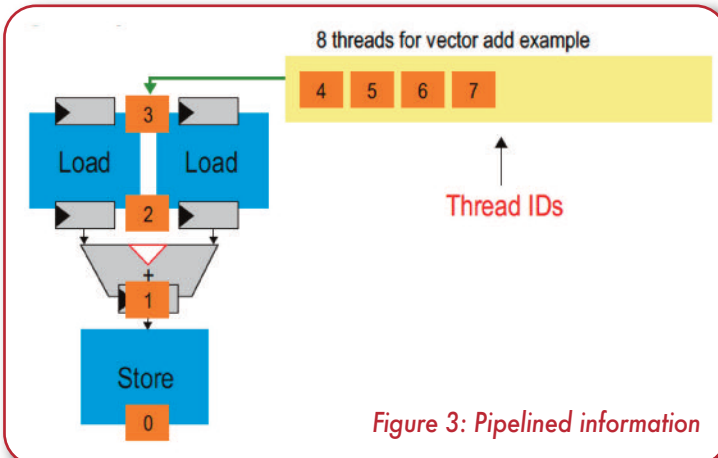

Figure 3: Pipelined information

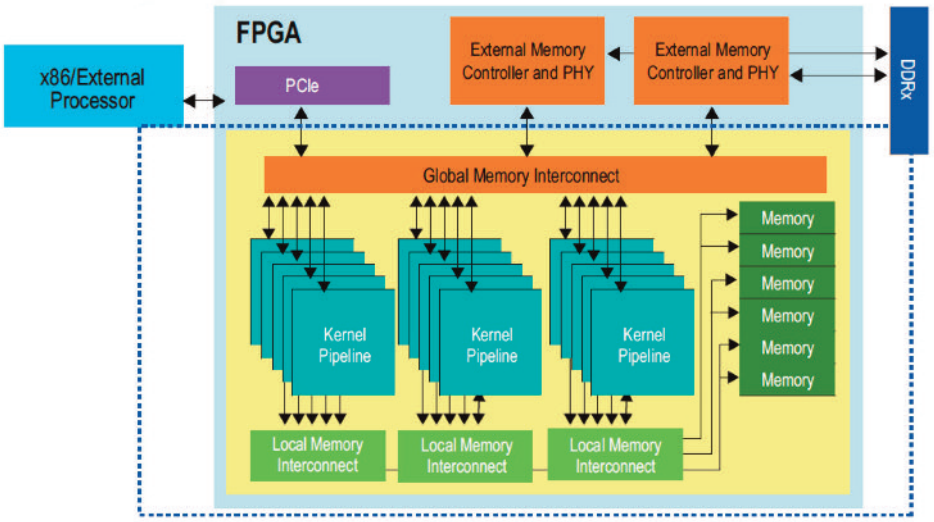

Figure 4: OpenCL system implementation

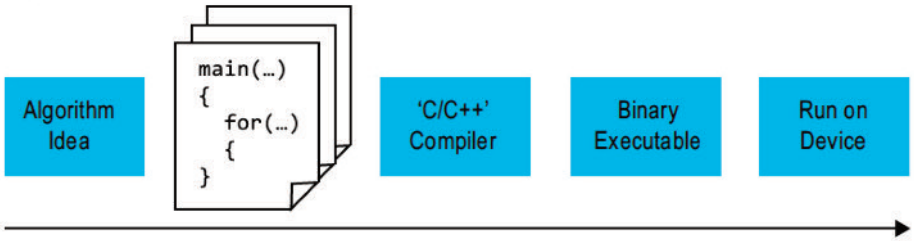

Figure 5: Programmer's view

that they should begin fetching the first elements of data from arrays A and B. On the second clock cycle, thread 1 is clocked in at the same time that thread 0 has completed its read from memory and stored the results in the registers following the load units. On cycle 3, thread 2 is clocked in, thread 1 captures its returned data, and thread 0 stores the sum of the two values that it loaded. It is evident that in the steady state, all parts of the pipeline are active, with each stage processing a different thread.

Figure 4 shows a high level representation of a complete OpenCL system containing multiple kernel pipelines and circuitry connecting these pipelines to off-chip data interfaces. In addition to the kernel pipeline, Altera's OpenCL compiler creates interfaces to external and internal memory. The load and store units for each pipeline are connected to external memory via a global interconnect structure that arbitrates multiple requests to a group of DDR DIMMs. Similarly, OpenCL local memory accesses are connected through a specialised interconnect structure to on-chip M9K RAMs. These specialised interconnect structures are designed to ensure high operating frequency and efficient organisation of requests to memory.

The creation of designs for FPGAs using an OpenCL description offers several advantages in comparison to traditional methodologies based on HDL design, the most significant of these is shown in Figure 5. Development for software-programmable devices typically follows the flow of conceiving an idea, coding the algorithm in a high-level language such as C, and then using an automatic compiler to create the instruction stream.

This approach can be contrasted with traditional FPGA-based design methodologies. Here, much of the burden is placed on the designer to create cycle-by-cycle descriptions of hardware that are used to implement their algorithm. The traditional flow involves the creation of datapaths, state machines to control those datapaths, connecting to low-level IP cores using system level tools (e.g., SOPC Builder, Platform Studio), and handling the timing closure problems since external interfaces impose fixed constraints that must be met.

The goal of an OpenCL compiler is to perform all of the above steps automatically for designers, allowing them to focus on defining their algorithm rather than focusing on the tedious details of hardware design. Designing in this way allows the designer to easily migrate to new FPGAs that offer better performance and higher capacities because the OpenCL compiler will transform the same high-level description into pipelines that take advantage of the new FPGAs.

Author profile: Deshanand Singh is the Supervising Principal Engineer, Software and IP Engineering, at Altera Corporation

More from Altera
Return to contents page.

Electronic Specifier design

Enabling a smarter life

How broader access to personalised information could promote a greater sense of community

Recent advances in information technology have brought many digitalised applications into our lives, such as automated Interactive Voice Response (IVR) over the phone, Skype, and Fax-on-Demand, making our lives more convenient than ever. Many of these applications have broken the boundary between analog data and digital data, and the newest developments are toward incorporating various forms of communications into a unified, PC-based messaging platform; including phone calls, e-mails, text messaging and voice mail, allowing users to receive information anytime, anywhere.

When this technological trend is applied to residence or community management, much of the household information can be digitalised and connected with the community information system, which will use a unified messaging platform to send messages to our computer, telephone, TV, smartphone, or other end user devices, allowing us to have firsthand information about our community, deepening our participation in community affairs and enhancing our community awareness. This is what we call 'e-community' and 'e-life'.

What are the differences between an 'e-community' and a 'non e-community'? For example, when a suspicious event happens in a community without informationised management, concerned residents can only visit the community office to review video surveillance records to find clues related to that event.

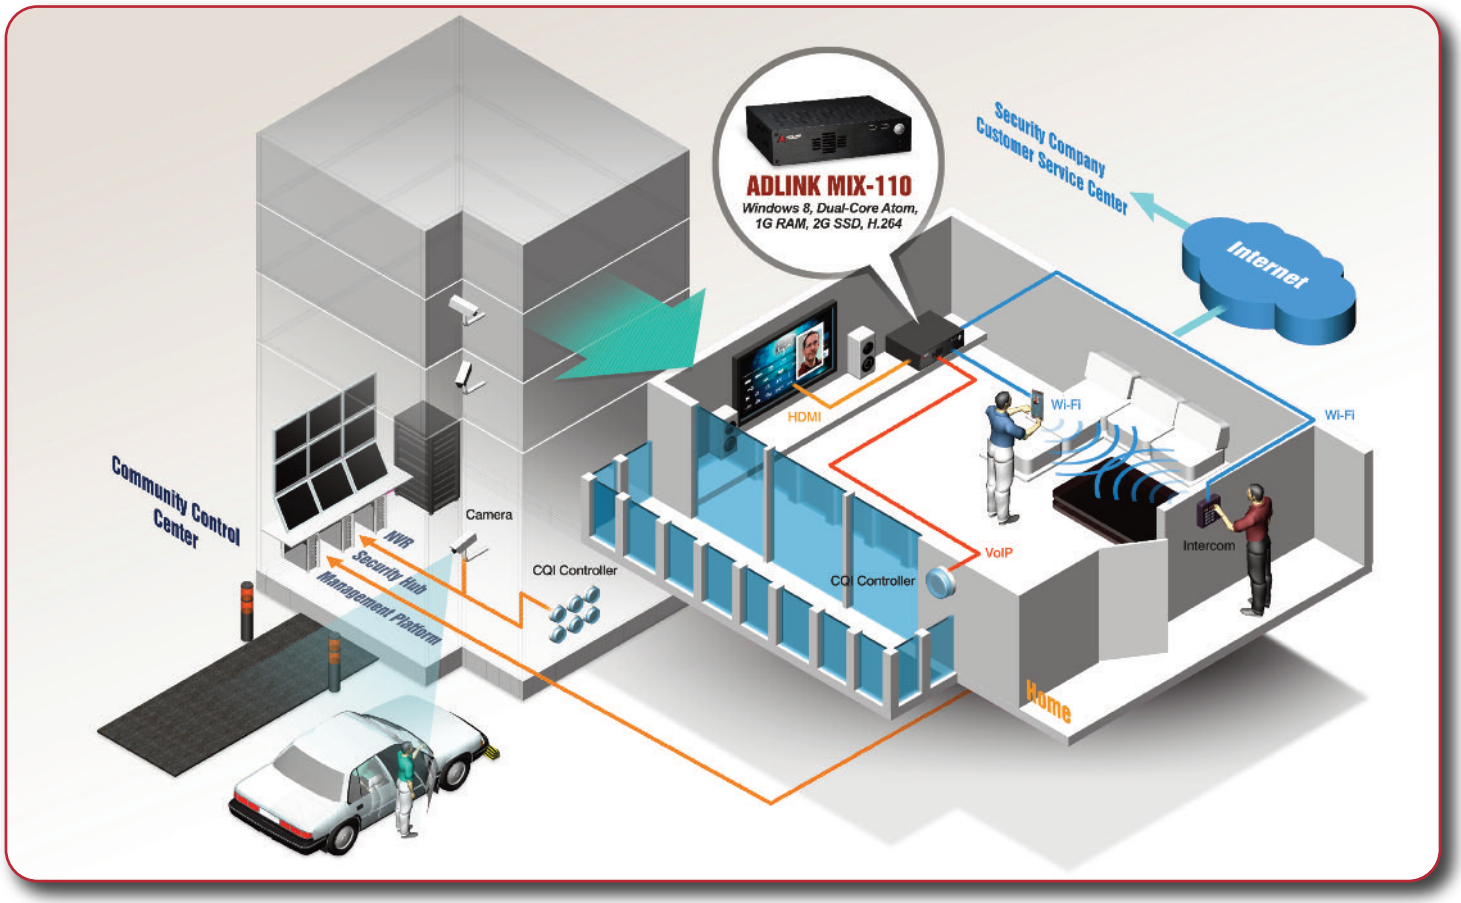

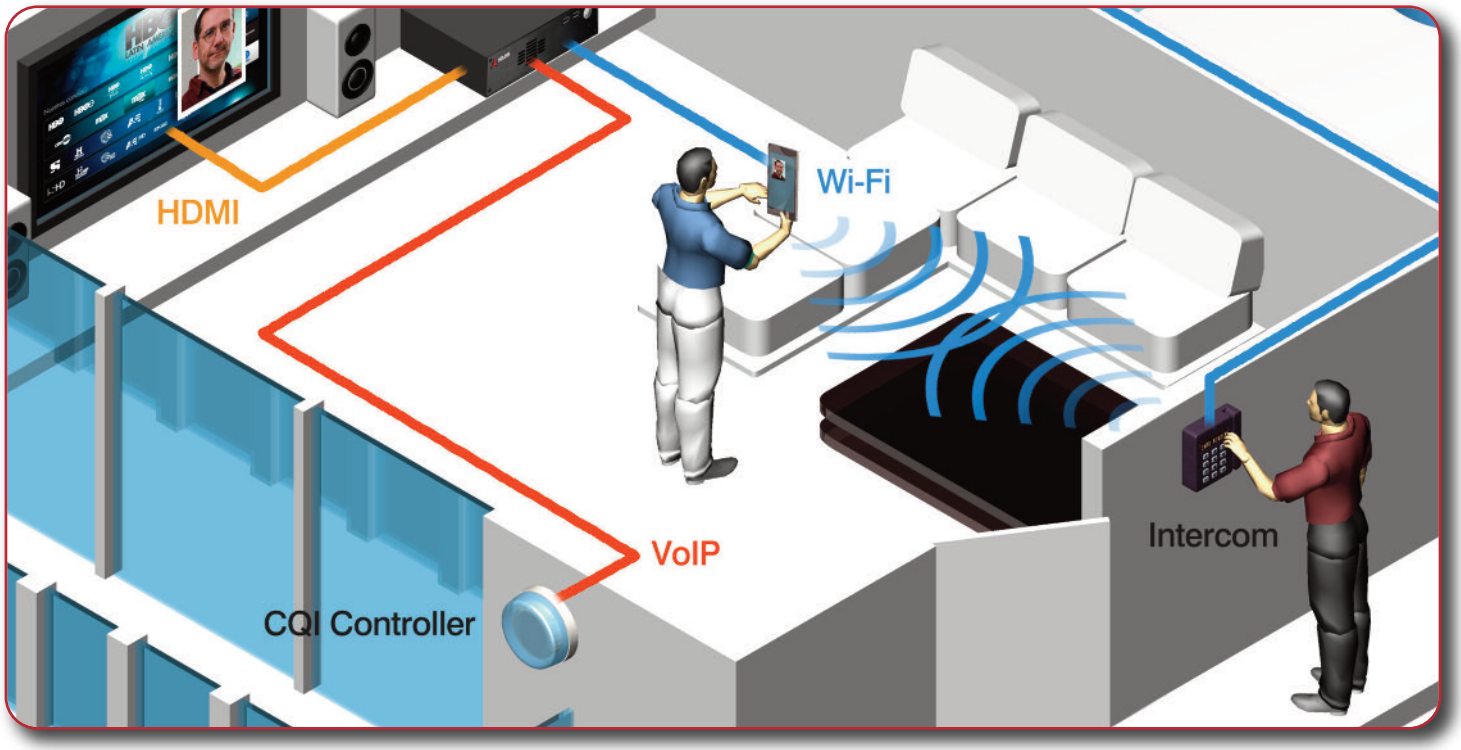

In contrast, residents in an e-community can view real-time and historical videos of their community on their TVs in the comfort of their homes, or by logging into the community website from a computer, smartphone or any other portables device wherever they are.

Living in an e-community, residents can also query and browse other information about their community via the Internet, for example, searching for upcoming events or checking the power usage of their community.

With informationised community management, the community management committee or an outsourced company can send notices and messages to residents via a computerised community management system. The management system can also be linked to a contracted security company to provide more services.

Such digital trends are making the management of a community more convenient and efficient and improving the well-being of community residents.

Application architecture

The deployment of an e-community requires a reliable computer system with rich interfaces for connecting to the community video surveillance system, entrance control system, intelligent power metering system, community management system, and security company customer service system.

This computer system serves as a computing platform to acquire data from connected devices, process the acquired data, store data, and display information as required by applications.

A commercial-grade computer is inadequate for carrying out such missions as they usually have a limited number of interfaces and are less stable and reliable. The computer system displays information on a display or TV screen in the community office, or sends required information to residents' terminal devices such as desktops or smartphones.

With the e-community system, residents can have firsthand information about their residence and their community when they are at home, at the workplace, or even when travelling thousand miles away from home.

Functions

Residents will be able to identify and talk to visitors ringing the doorbell when they are in a bedroom, living room, kitchen or even bathroom simply with a hand-held device (tablet or smartphone), through the TV, or with other devices which have HDMI or Wi-Fi interfaces, without having to move to the intercom to answer it.

Residents can review the power consumption status of their household and of community

public areas through handheld devices or computers, so that they can manage their power use behaviour.

Video cameras are installed throughout the public areas in the community and linked to a control centre. Security guards can review real time video records in the control centre to oversee the security situation in the community, and residents can use their own TV or hand-held devices to watch their children playing in the community as well as vehicles entering or leaving the community.

Residents can receive community information via various kinds of devices. A contracted security company can receive real time updates on security situations in the community, and dispatch persons to handle emergency situations if necessary.

Residents can set links to contracted companies and allow them to get specific data about the health condition of their family members or the operation status of their electronic appliances so that those companies can provide services if they detect any abnormal situations. Meanwhile, with the use of video cameras, residents can observe the status in their home with mobile devices when they are away.

The ADLINK MIX-110 is a small form factor (234.4mm x 191.2mm x 56.6mm) embedded system with a compact design suitable for a community management system. Running on a power-saving Intel Atom Processor D2550, the MIX-110 provides up to 6 USB ports, 2 serial COM ports, 1 Mini PCIe expansion slot, one general purpose IO and 2 SATA 3Gb/s slots. These interfaces can be used to connect with the community's network video recorders (NVR), CQI smart power meters, RFID-enabled entrance control card readers, community management system as well as a security company's customer service centre network.

The MIX-110 provides dual display with VGA and DVD-D outputs to support high-definition screens set up in the community committee office. Meanwhile, wired and wireless communications are available for connecting with terminal devices in the homes of community residents.

Featuring rich I/O, small dimensions, low power consumption and ADLINK's reliable design, the MIX-110 delivers stability and reliability, and is an ideal choice for informationised community management.

More from Adlink Technology
Return to contents page.

Under inspection

Tackling rising network traffic with high-performance deep packet inspection for increased security and intelligent packet handling. By Steve Konish

The network infrastructure is becoming overwhelmed by ever-increasing data traffic, resulting in a growing number of security breaches and attacks of malicious content. Equipment providers aiming to deliver increasingly intelligent solutions for mobile, fixed and enterprise networks are searching for new software-enabled deep packet inspection (DPI) solutions such as pattern matching and content scanning to meet these challenges. Using high-performance multi-core technologies, software-enabled DPI can also deliver significant advantages in the area of application identification, traffic shaping and policy server implementations, enabling providers to better manage networks as well as being able to add new services, including those at tiered levels.

Billions of devices and machines are now connected to the network. This number is expected to triple in the coming few years as connectivity continues to rapidly expand beyond PC, smartphones and tablets and deeper into the world of machine-connected devices of the 'Internet of Things'. These devices and machines will generate massive amounts of new data traffic on the network. The impact on the existing infrastructure will create serious challenges for equipment providers and network operators – in fact, the explosion of connected devices will generate a level of traffic that has the potential to bring existing networks to their knees. As the demand for network usage increases, new services will be needed to deliver higher value, greater security and better quality of experience. Meeting these challenges will require new and more intelligent solutions, driven by advances in software that take full advantage of new processing technology and effectively employ technologies such as DPI.

Use cases for DPI

There are two primary use cases for DPI: firstly using the technology for various different kinds of security applications; and secondly using it in policy servers and flow analysis to analyse and direct traffic and bandwidth. And for these use cases there are three key target markets, which are: Communications systems of all kinds; enterprise networks, which could be any kind of campus network or metropolitan area network (MAN); and wireless networks, such as 3G, LTE, radio-access and mobile broadband networks.

While these target markets utilise different network elements, all still have requirements for both security and policy server/flow management. As an example, a security DPI solution is critical to the enterprise network with its various requirements for firewalls, threat management, virus protection, and intrusion detection and protection. This is also the case for radio networks, but only in certain places at the access layer, such as where the radio-network front-end is accessible in a particular area or cell. Mobile broadband also requires different types of devices that demand both security and policy management.

According to the latest report from Infonetics Research, which tracks DPI software solutions and related hardware deployed in wireless and fixed-line networks, service provider DPI revenue was up by 28% in 2012, totalling $596 million. While fixed-line operators are still a significant portion of the DPI market for data traffic management, it is mobile that is expected to drive the DPI market over the next several years. In the same report, the wireless segment of the DPI market was forecast by Infonetics to grow at a 33% CAGR from 2012 to 2017.

Certainly, a fast growing trend in the industry is the ability for network operators and Internet

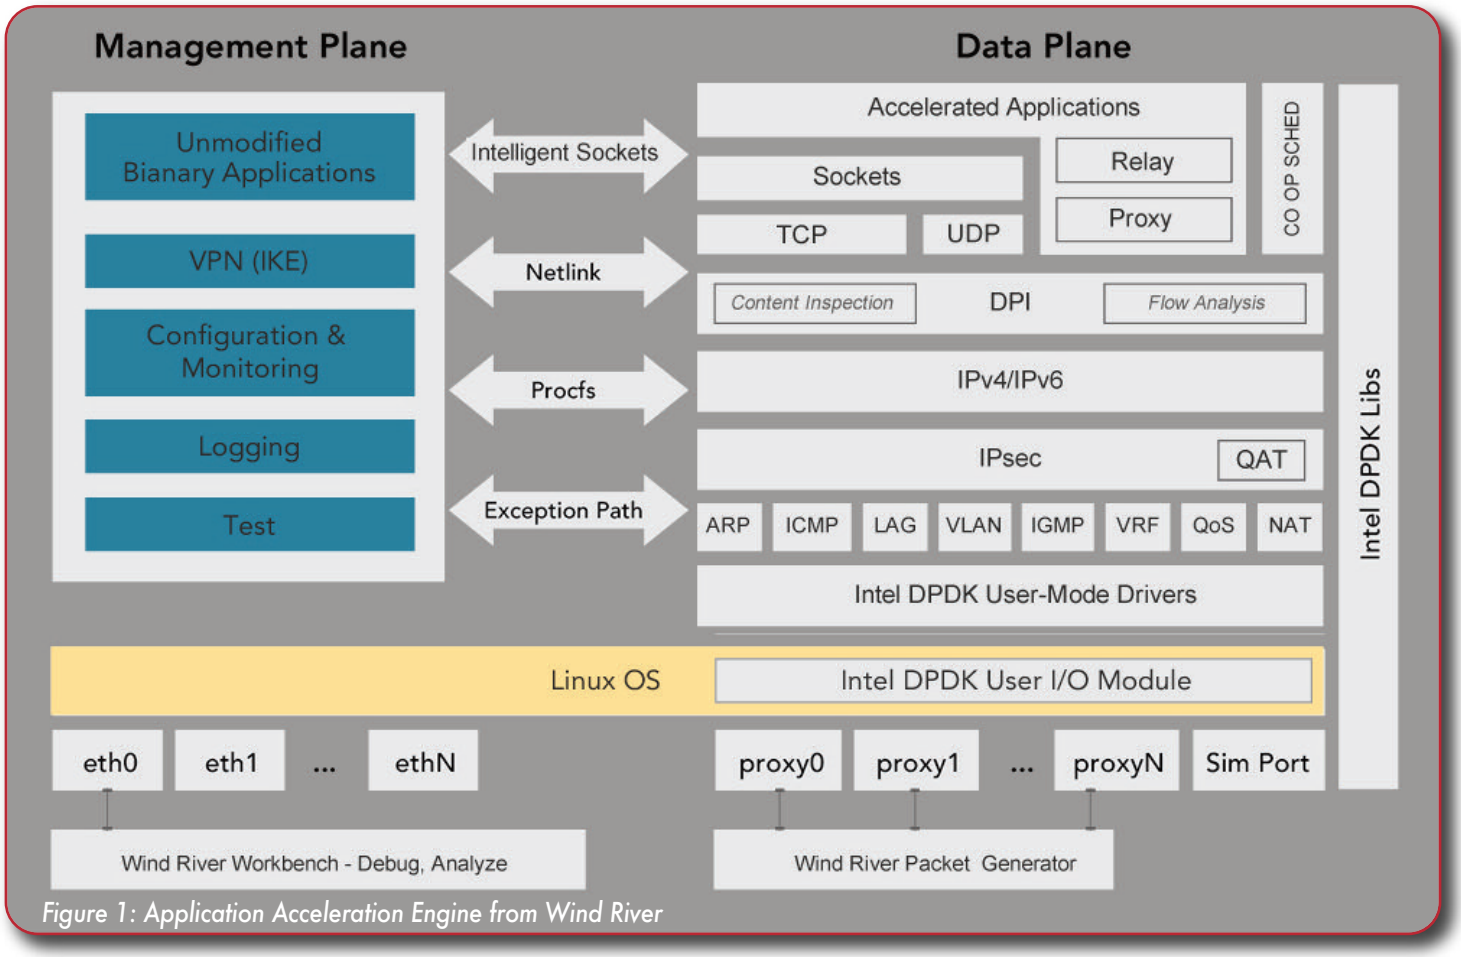

Figure 1: Application Acceleration Engine from Wind River

providers to offer differentiated levels of service. QoS (Quality of Service) policy filters that employ DPI technology to determine the origin of the packets are becoming increasingly pervasive, enabling improved levels of subscriber-aware traffic management. During peak usage times, for example, the technology could enable providers to give a lower priority to data that is associated with only a certain level of subscriber service, perhaps for high-bandwidth applications such as video streaming.

Deep inspection

As the names suggests, DPI analyses entire packets and correlates information across multiple packets to identify the application, whereas 'shallow packet inspection' only analyses the header of a data packet. The traditional security firewall for example will look only at packet headers. Advanced security application equipment will use DPI technology to examine the content of a data stream to detect threats rather than basing security decisions solely on fields in the header. The detailed examination of packets has an added cost, however, because scanning the content of data streams is CPU intensive and can become unfeasible as the load increases. The effectiveness of DPI therefore depends largely on how well it performs under heavy traffic.

To prevent this from occurring and to increase performance, network security vendors typically relied on hardware-assisted DPI technology to perform detailed content scanning at network speeds, using expensive purpose-built silicon such as application-specific processors, custom ASICs or FPGAs. In addition, this hardware has not offered the high levels of flexibility needed to adapt to the changing needs of networks, as well as being a solution that often leads to significant development and manufacturing costs.

However, DPI has been undergoing a rapid evolution in the market with the availability of advanced and high-performance multi-core processors from silicon vendors such as Intel. This has created the opportunity for well-designed DPI software to approach and even exceed the performance of previous hardware-based solutions. A software-based approach to intelligent and high-speed DPI can therefore provide a cost-effective and scalable solution that

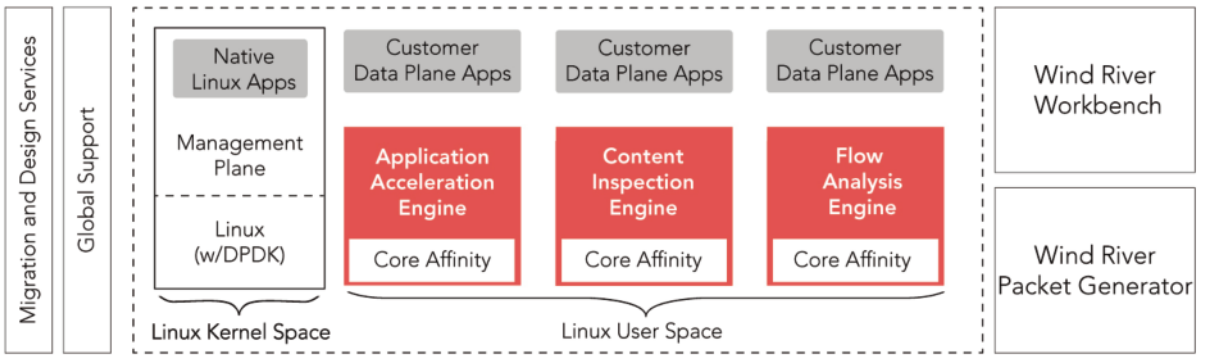

Figure 2: Key components of the Intelligent Network Platform

also has the flexibility to evolve with changing system requirements, especially in the security area where new malware and viruses are occurring constantly.

Intelligent pattern matching

Central to many DPI implementations is pattern matching: the ability to compare and match incoming data streams with a database of known offending patterns called signatures. These signatures represent potential malicious content and can take the form of simple literal strings or more complex patterns such as a specific arrangement of bytes broken up by a variable amount of other possibly irrelevant data.

While literal searches can be quite intensive, regular expression searches require significant processor resources that can become problematic to perform at high speeds, especially when searching for thousands of signatures. An example of a leading software pattern-matching solution that compares large groups of regular expressions against blocks or streams of data is the Wind River Content Inspection Engine (CIE). Designed for low-end single-core to high-end multi-core processors, it is a scalable and cost-effective DPI approach that runs entirely in software. CIE can be scaled from less than 1Gbps up to 160Gbps, depending on the number of cores being used. While the engine has been designed to take full advantage of Intel multi-core processor technology, it is also processor architecture and operating system independent, therefore providing the opportunity to upgrade units already installed in the field. Additionally, easy to upgrade software updates provide continued performance improvements and functionality.

Ideal for applications that need to scan large amounts of data at line rate, such as intrusion prevention (IPS), antivirus (AV) and unified threat management (UTM), among other DPI systems, it delivers a cost-effective software solution for scanning data content at line rate in security equipment ranging from small network appliances to large network elements.

Application acceleration

As well as the issue of security, the growing number of connected devices means that network operators also face significant challenges to increase performance and find efficient and scalable ways to handle the on-going explosion in network traffic. This is in addition to expanding their network services to achieve higher average revenue per user (ARPU). Performance, however, remains a vital priority. Building in greater intelligence with DPI is highly desirable, but operators will not want to sacrifice speed for it.

Wind River Application Acceleration Engine is a comprehensive, optimised network stack designed for the acceleration of ISO-model layer 3 and 4 network protocols and the applications that run above them. Delivering major performance boosts compared to the native Linux network stack, the Application Acceleration Engine can provide up to a 1100% improvement in IP-forwarding for select packet sizes, up to a 500% improvement in throughput for User Datagram Protocol applications and a performance boost for TCP. Use of this engine means that network applications of all kinds, both legacy and next generation, can significantly improve performance and enable a number of new, high-value and extremely intelligent network services.

The Application Acceleration Engine works in the data plane in conjunction with Linux. While not exclusively, the software engine has been fine-tuned for Intel multi-core processors and

communications platforms. Specifically the engine builds on the performance of Intel Data Plane Development Kit (DPDK) to accelerate networking applications of all kinds, and security components such as DPI and virtual private network (VPN) technologies. The Intel Data Plane Developer Kit (DPDK) is a set of highly optimised user-space libraries and drivers that enable the consolidation of control and data plane platforms and execute efficient packet processing on Intel architecture (IA) processors.

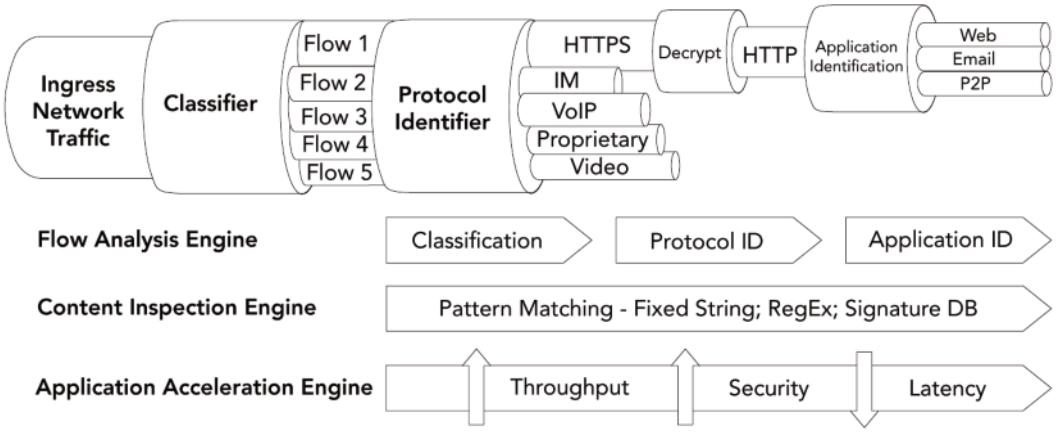

Figure 3: Network acceleration, deep packet inspection and packet identification in one system

The Application Acceleration Engine can effectively isolate optimised application protocols including TCP and UDP and IPsec on a processor core and execute them in non-blocking mode without the need for context switching via the Linux operating system or IP stack. Delivering almost hardware-level throughput, it essentially takes the packet from memory or from the Ethernet port straight up to the DPI application through the optimised protocol stack, meaning the application can process the packet in almost real time. The Application Acceleration Engine is shown in Figure 1.

The Wind River Intelligent Network Platform integrates both the Content Inspection Engine and the Accelerated Application Engine along with other key software components to allow the design of high-performance intelligent network applications in a consolidated management and data plane system. The Intelligent Network Platform can be used in its entirety or in a modular fashion, in many cases, with companies interchanging and extending individual components with their existing in-house technologies. Key components of the platform are shown in Figures 2 and 3.

In addition to the Content Inspection Engine and Application Acceleration Engine, the complete platform also includes a Flow Analysis Engine, which offers complete visibility into network traffic in real time, including flow classification, protocol and application identification, and metadata extraction. Also included is the Intel Data Plane Development Kit (Intel DPDK) library, plus optimisations for carrier grade Wind River Linux, including the flexibility to support other Linux distributions. The platform enables the acceleration of Linux-based applications up to three times faster without the need to modify code in existing applications. Even greater performance can be achieved of up to five times faster if an application is specifically ported and optimised for the platform.

Key applications for the platform include: cloud radio-access networks; intrusion detection and prevention applications, media gateways; mobile access, edge and core; next-generation firewalls; server acceleration and offload; and virtual private networks (VPNs).

In summary, Wind River Intelligent Network Platform is a software platform for the development of sophisticated network equipment that can accelerate and secure the flood of traffic for current and future networks. Additionally, the high-performance of multi-core processing technology from Intel and others have helped realise the possibilities for service providers and network equipment suppliers to use leading-edge software-based DPI solutions to accelerate, analyse and secure applications, in turn enabling increasingly intelligent solutions for mobile, fixed and enterprise networks.

Author profile: Steve Konish is the Director of Product Management, Communication Platforms, with Wind River

More from Wind River
Return to contents page.

Electronic Specifier

The engineer's resource guide to new products, technology and design

For full details on any of the Electronic Specifier products
please email our Advertising Sales Manager, Ben Price ben.price@electronicspecifier.com

www.electronicspecifier.com